

UC Berkeley

UC Berkeley Previously Published Works

Title

Foundation models for atomistic simulation of chemistry and materials

Permalink

<https://escholarship.org/uc/item/2rq401k4>

Journal

Nature Reviews Chemistry, 10(3)

ISSN

2397-3358

Authors

Yuan, Eric C-Y

Liu, Yunsheng

Chen, Junmin

et al.

Publication Date

2026-03-01

DOI

10.1038/s41570-025-00793-5

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Foundation Models for Atomistic Simulation of Chemistry and Materials

Eric C.-Y. Yuan^{1,6}, Yunsheng Liu¹, Junmin Chen¹, Peichen Zhong⁵, Sanjeev Raja^{4,7}, Tobias Kreiman⁴, Santiago Vargas^{6,8}, Wenbin Xu⁹, Martin Head-Gordon^{1,6}, and Chao Yang⁷

Samuel M. Blau^{*8}, Bingqing Cheng^{*1,6}, Aditi Krishnapriyan^{*3,4,7}, Teresa Head-Gordon^{*1-3,6}

¹Kenneth S. Pitzer Theory Center and Department of Chemistry, ²Department of Bioengineering, ³Department of Chemical and Biomolecular Engineering, ⁴Electrical Engineering and Computer Science, ⁵Baker Institute of Digital Materials for the Planet, University of California, Berkeley, CA, 94720 USA

⁶Chemical Sciences Division, ⁷Applied Mathematics and Computational Research Division, ⁸Energy Technologies Area, ⁹National Energy Research Scientific Computing Center, Lawrence Berkeley National Laboratory, Berkeley, CA, 94720 USA

*corresponding authors: aditik1@berkeley.edu, bingqingcheng@berkeley.edu, smblau@lbl.gov, thg@berkeley.edu

Abstract

Given the power of large language and large vision models, it is of profound and fundamental interest to ask if a foundational model based on data and parameter scaling laws and pre-training strategies is possible for learned simulations of chemistry and materials. The scaling of large and diverse datasets and highly expressive architectures for chemical and materials sciences should result in a foundation model that is more efficient and broadly transferable, robust to out-of-distribution challenges, and easily fine-tuned to a variety of downstream observables, when compared to specific training from scratch on targeted applications in atomistic simulation. In this Perspective we aim to cover the rapidly advancing field of machine learned interatomic potentials (MLIP), and to illustrate a path to create chemistry and materials MLIP foundation models at larger scale.

1 Introduction

The fundamental principle underlying the description of atomistic chemistry and materials science is the Schrödinger equation, whereby a chemical system and its physical properties can be uniquely defined by Cartesian and spin coordinates of the component atoms and system net charge. But as famously stated by Dirac after the discovery of quantum mechanics "...the underlying physical laws necessary for the whole of chemistry are thus completely known, and the difficulty lies only in the fact that the exact application of these laws leads to equations much too complicated to be soluble." This has led to a series of well-controlled approximations, such as the Born-Oppenheimer separation of nuclei and electronic degrees of freedom, and the transformation of the Schrödinger eigenvalue equation into an algebraic framework, that has allowed for good model chemistries to be solved for systems beyond the hydrogen molecule.¹ Just as important is the need to satisfy the laws of statistical mechanics that requires the high dimensional integration over coordinates of Boltzmann weighted configurations to predict observables, usually collected over time trajectories with the hope that ergodicity is satisfied. Even so, the most widely used approximate electronic structure methods based on Density Functional Theory (DFT)² combined with ab initio molecular dynamics (AIMD) remains inaccessible for long timescales and large length scales and system sizes that define many of the interesting areas of chemistry and materials science.

How might we resolve this tension between the need for quantum mechanical accuracy with the need to satisfy statistical mechanical sampling to yield converged and correct observables? One recent possibility is to leverage the tools of machine learning and data science. Their success in other domains suggests that we could learn "the foundations" of molecular and materials chemistry from an abundance of data using appropriate architectures and training strategies. The concept of a foundation model (FM) has been established by the natural language processing and computer vision communities through the development of large language models (LLMs) and large vision models (LVMs).^{3,4} These large parameter architectures are pre-trained in an unsupervised or self-supervised manner on enormous amounts of data, and can then be easily fine-tuned for accurate prediction on seemingly unrelated downstream tasks, revolutionizing the field of artificial intelligence.

While LLMs and LVMs themselves hold great potential to impact chemistry and materials science,⁵⁻⁹ in this perspective we are focused on whether FMs can directly capture the potential energy surface (PES) via machine learned interatomic potentials (MLIPs).^{10,11} The end goal of such a model is to enable large-scale atomistic simulation and analysis, including the ability to derive any observable of interest, for any molecular or periodic system, either with a small amount of system-specific fine-tuning¹² or without any fine-tuning at all. Importantly, it would supersede the painful acquisition of new labeled data and the hidden expense of often hundreds of rounds of retraining of the ML model used in active learning to develop a robust MLIP.¹³⁻¹⁶

To achieve such a goal, this perspective first provides more concrete definitions for what defines a FM, and how a FM for atomistic simulation should be distinguished from a "universal potential" or even a more restricted definition such as transfer learning. In particular, our criteria begins with the ability of a MLIP FM to exhibit scaling laws and pre-training and data strategies that have made LLMs and LVMs so powerful. With this context, we examine the capacity of current model architectures and identify limitations and opportunities that

are relevant to their scalability to larger learnable parameters in order to become more expressive and foundational to increasing amounts of data. We also consider what type and amount of data of scale will be required for a generally applicable pre-trained FM, how that compares to the data that already exists, and where we expect fine-tuning data to be most important for chemistry and materials. In this perspective we focus primarily on FMs trained on energy and force data to address outstanding problems in chemistry and materials through molecular simulation. But it should be pointed out at this juncture that different modality data exists, such as images and text extraction for catalytic discovery,¹⁷ prompt engineering of LLMs for predicting synthetic pathways of small drug molecules,⁹ and use of video data such as that generated for liquid-phase TEM imaging.¹⁸ We further examine the importance of more advanced training strategies, including self-supervised and unsupervised pre-training as well as model distillation, and highlight that new methods for physical infusion need to be relevant and feasible when training at scale on advanced hardware. Finally, we consider how the performance of MLIPs is currently evaluated, and to suggest new benchmarks for demonstrating scaling laws and the ability to be broadly fine-tuned for a wide range of downstream applications, which is particularly important given the breadth of desired observables in chemistry and materials.

2 Characteristics of Foundation Models

The term “foundation model” has been loosely defined to mean any large-scale model pre-trained on large amounts of diverse data to capture a broad range of complex patterns,⁴ as represented by popular LLMs such as GPT¹⁹ and Llama.²⁰ One of their defining characteristics is that they often obey heuristic scaling laws,^{3,21–23} which refers to how their performance improves as a function of increasing model size, training data, and computational resources. Increasing model capacity improves expressivity and performance when coupled with improved algorithms for neural network operations, such as attention mechanisms most commonly used in highly scalable Transformer²⁴ architectures. Data scaling for pre-training refers to increasingly large and diverse datasets of typically unlabeled data to learn general features and patterns, covering various domains and contexts to ensure broad knowledge acquisition. Compute scaling permits parameter and data scaling through optimal use of GPU and CPU hardware, exploiting single and mixed precision computation when appropriate, and taking advantage of parallelization using distributed computing. Exhibiting and satisfying these scaling laws has been critical for FMs in achieving strong results across various applications and domains, ranging from natural language processing, computer vision, climate modeling, and robotics.^{25–28}

Coupled with the scaling laws is a pre-training strategy that ultimately yields a high efficiency model that performs well across a series of benchmark data sets that exist for fine-tuning for down-stream tasks. Pre-training objectives, particularly in language, often involve self-supervised learning tasks such as predicting missing words in sentences or predicting the next word in a sequence. During the earlier stages of the training process, more compute is used while data quality is typically lower and human supervision is minimal. The resulting pre-trained scaling models, data, and compute can lead to emergent capabilities, where a FM becomes capable of solving a task that appeared not possible at smaller scales.²⁹

Although it is still disputed how and why capabilities are unlocked at scale,³⁰ these laws provide a quantitative framework for understanding and predicting the behavior of chemical FMs as they are scaled up, and offer insights into how different factors (architecture, data, pre-training) contribute to model performance improvements.

One important clarification that we hope to address in this perspective is to distinguish the term "foundation model" from other related terms such as "transferable" or "universal", or large atomic models (LAMs),³¹ particularly in the context of MLIPs. The popular MLIP MACE-MP-0³² is trained on data from the Materials Project^{33,34} generated at the PBE level of DFT. The resulting ML model exhibits good accuracy in energy and force predictions against PBE on other crystalline unit cells and exhibits excellent stability in the numerical solution to Hamilton’s equation of motion across related chemical and materials systems. This has opened up a whole new range of applications, including free energy calculations, simulating spectroscopic observables, and explorations of phase diagrams. However, we would stipulate that MACE-MP-0 is not a foundation model, but an excellent universal PBE potential, given that it is trained with supervision to do a very specific task: predict energy and force labels for a single chemical domain at one level of theory. Another key limitation is that MACE-MP-0,³² like many universal MLIPs, does not attempt to distinguish between systems with different total charge and spin, whose energy and forces would be different although the underlying input structure is the same.

We offer that to be a true FM preferably all of the following features should be demonstrated. First is that the zero-shot or fine-tuned MLIP FM should show performance across a broad range of down-stream tasks that is superior to task-specific models trained from-scratch.^{31,35} Second, MLIP FMs would show compliance with heuristic scaling laws in regards how performance improves as model parameters and training data increases with increasing compute resources. Finally, a demonstration that a large-scale FM MLIP would have emergent capabilities, for example predicting higher quality CCSD(T) data or magnetic field-dependent properties, properties that are very different than the original underlying Cartesian geometries and DFT energy training data.

3 Foundation Models for Atomistic Simulations

3.1 Machine Learning Architectures for Chemistry and Materials

Over the past two decades, MLIP architectures have evolved significantly, integrating the inductive biases of chemistry while incorporating advancements from other areas of machine learning. We provide a general survey of these developments, as well as an outlook on recent architectures which increasingly exploit scale as a driving factor in performance. The architectures of MLIPs are posited to be distinct from those developed in computer science for tasks like natural language processing, the reason being given that chemistry is governed by physical laws and constraints that are more inviolable than the statistical nature of probability distributions in language. Hence most MLIP architectures are typically "physics-informed", while the exact manner in how to realize the physical constraints has gone through a few waves of evolution.

The early Behler-Parrinello neural network potentials³⁶ used hand-crafted two-body and three-body descriptors for describing atomic environments. Later, DeePMD models³⁷ automated the optimization and discovery of such descriptors, reducing reliance on manual feature engineering. The atomic cluster expansion (ACE) method³⁸ introduced a unified and generalizable framework for constructing atom-centered descriptors using systematically improvable body-ordered terms. Graph neural networks (GNNs) then broadened the landscape, allowing for flexible and expressive architectures that learn features by iteratively exchanging information between neighboring atoms, thus capturing many-body effects.^{39–42} More recently, GNN models have incorporated different forms of graph attention mechanisms.^{43,44} These models can be viewed as a more general way of performing message passing between atoms.

Despite the rapid emergence of diverse architectures, most MLIPs incorporate similar inductive biases. The first is the nearsightedness assumption. This postulates that energy and forces experienced by a central atom are primarily influenced by its neighboring atoms within a finite cutoff radius. Nearsightedness readily ensures the key model chemistry requirement of size-consistency¹ is satisfied. Message passing schemes relax this assumption to some extent by enabling atoms to aggregate information from progressively distant neighbors through multiple iterations, although intermediate neighbors still contribute the most. Second are the physical constraints such as energy conservation, smoothness, and invariances (translational, rotational, and permutational). These constraints are incorporated either explicitly within the model architecture (e.g., through gradient-based forces, symmetrized input features or invariance-preserving transformations) or implicitly via the loss function and training data. The third is to describe interatomic interactions via body-order terms. Two-body (pairwise) terms dominate, while three-body terms are also critical. Higher-body terms exhibit diminishing returns due to increased computational cost and reduced relative contributions. Frameworks like the Atomic Cluster Expansion (ACE)³⁸ provide systematic control over the inclusion of higher-body terms, with user-specified truncation. In GNNs, adding a message-passing layer effectively increases the body order by one to yield models of different ranks.

Moreover, the seemingly diverse existing architectures share the same mathematical foundation. GNNs are all based on the same representations where atoms serve as nodes and their distances within the cutoff radius are edges. Additionally, almost all atomic and structural representations for materials and molecules are fundamentally related and can be understood within the unifying framework of atomic density^{45,46} or, equivalently, ACE.³⁸ Moreover, these representations can also be viewed as specific graph features, encoding geometric and chemical information of the nodes. This duality highlights a deeper connection: most existing MLIPs can be understood as different variations of the same underlying framework.

The unifying framework of the current MLIP architectures and the common approach to enforce physical constraints are to be contrasted with the well-known “bitter lesson” that argues that general methods that leverage computation and data are more effective in the long term over approaches that try to build knowledge into the ML pipeline.⁴⁷ The main limitations of constraints are the challenges associated with the training dynamics and resulting poor loss landscapes, as well as difficulty with parallelization and distributed training on modern GPU hardware.^{48–50} Indeed, we are now seeing that scalable architecture methods that emphasize expressive capacity can both outperform models with built-in constraints

and effectively learn these constraints directly from the data, while being significantly more compute-efficient.^{27,28,44,51-53} This offers the first evidence of the important bitter lesson learned from other fields of ML that also have constraints, in that fewer inductive biases do seem to eventually work better. It further suggests that satisfying data and parameter scaling for chemistry and materials will be more effective compared to models that build constraints into the model.

One of the most well-studied physical constraints is rotational equivariance.^{41,42,54} It is worth noting that rotational equivariance is an easy property to satisfy, and cheap data augmentation (rotational transformations of the existing data) is more than adequate in nearly all cases.^{44,45,55} It is a more open question whether more complicated physical laws can be learned from large data or still require some inductive biases. An example would be whether the symplectic structure of Hamilton’s equation of motion can be respected, i.e. to drive stable molecular dynamics (MD) simulations. Stability in MD requires two key conditions: the first is the absence of pathological behaviors, such as molecules "exploding" due to data insufficiencies and numerical instabilities, even at very small time discretizations (Fig.1e). Second, is the ability to perform MD simulations without introducing non-conservative and energy drift artifacts. Recent studies have observed that certain MLIPs that use directly predicted non-conservative forces rather than obeying the energy conservation law can lead to issues such as artificial heat generation in NVE ensembles or severely modified dynamics in NVT ensembles.⁵⁶ These behaviors undermine the reliability of such non-conservative MLIPs for long-timescale dynamics and their integration into predictive downstream workflows. We also note that energy conservation is a constraint that can be agnostic to the architecture itself, as the standard way it is enforced is by taking gradients of the predicted potential energy in the loss function,⁴² an expense that is avoided in recent direct force models.^{54,55}

There are possible issues of the current MLIP architectures that should be examined more closely. The first is transferability outside the pre-training data. While recently developed universal MLIPs demonstrate reasonable performance in predicting materials stability on element-wise out-of-distribution datasets,^{57,58} these evaluations primarily focus on local energy minimum searches through structure relaxation of crystalline materials, sharing similar prototypes in the Materials Project database. A systematic softening (underestimation) of the ML-predicted PESs and interatomic forces⁵⁹ compared to DFT ground truth was observed in a recent benchmark study as seen in Figure 1a.⁶⁰ The underestimation of the PES is particularly significant for high-energy atomic configurations, which limits its application to molecular dynamics simulations and rare-event samplings. Recent large data efforts^{31,61} and open-sourced universal MLIPs^{62,63} that include more off-equilibrium atomic configurations show promise in addressing these challenges as discussed in Section 3.2. However, direct transferability of universal MLIPs for specific downstream molecular or materials modeling remains uncertain, as they are not yet demonstrably better than models built directly for the new system or property from scratch. We return to this issue in Section 3.5.

Another limitation is the lack of incorporation of long-range interactions. Although short-range potentials may be sufficient to describe most properties of homogeneous bulk systems,⁶⁴ they may fail for interfaces,⁶⁵ dielectric response,^{66,67} and dilute ionic solutions with Debye-Hückel screening. For example, consider two charged molecules that are separated by a distance greater than the cutoff radius used in most MLIPs (Fig.1b). These models, which focus on short-range interactions, will fail to accurately describe the electrostatic forces

between the molecules (Fig.1c). This issue cannot be resolved simply by using message passing, as the molecules exist on separate graphs and do not interact within the framework of short-range biases inherent to message-passing neural networks. On the other hand, MLIPs with built-in long-range corrections^{68–70} can effectively solve such artifacts (Fig.1d).




Figure 1: (a) Distribution of softening scales ($f_{\text{MLIP}}/f_{\text{DFT}}$) sampled from 1,000 high-energy configurations using the structures from the WBM dataset.⁵⁷ Reproduced with permission from Ref.⁶⁰ (b) An example of two charged molecules that are far apart and do not interact due to (c) short-range implicit bias from GNNs and message passing, and (d) a new ML model that incorporates long-ranged information.⁶⁹ Blue points are for training and red points are for testing. Adapted with permission from Ref.⁶⁹ (e) Depiction of instability in molecular simulations with MLIPs, which can occur even at very small time discretizations.⁷¹ Adapted with permission from Ref.⁷¹

There is also the concern of scalability: whether scaling laws will be exhibited as the community acquires larger and larger data sets and computational resources. Is the GNN architecture adequate, or is there a need to shift towards more scalable architectures^{72,73}? Relatedly, computational inefficiency when applying MLIPs at scale⁷⁴ can limit the application to realistic materials and chemical problems. The current short-ranged MLIPs are already orders of magnitude more expensive than traditional force fields. Further incorporating long-range interactions may require more computational power, and while there is a willingness to pay this price for improved accuracy, the current MLIP frameworks are not fully optimized for such tasks. While GNNs provide a powerful framework for capturing atomic interactions, their computational inefficiency at scale necessitates innovative approaches that can maintain or even enhance their model performance.

3.2 Data Requirements, Pitfalls, and Challenges

Real-world image and language datasets often contain trillions of tokens, whereas simulated chemistry datasets for training MLIPs typically contain between 1 to 200 million molecule or material configurations with associated energy and force data. Thus, increases in dataset size remain critical for improving the pre-training of MLIPs if they are to ultimately serve as robust FMs. Table 1 summarizes molecular and materials datasets with over 1,000,000 energy- and force-labeled structures that have emerged over the last ~ 5 years. The underlying labeling methods are DFT calculations with variable basis set sizes. We also include large-scale non-DFT chemical data such as Uni-Mol2,⁷⁵ which is a curated subset of the Zinc20 molecular data set,⁷⁶ and more recently Zinc22⁷⁷ which is comprised of 4.5 billion molecules that also contains force field derived labels of their physical properties, including partial atomic charges, cLogP values, and solvation energies.⁷⁷ In all cases this variable quality in data is highly appropriate for pre-training where more value should be placed on broadly covering chemical space, with domain diversity i.e. molecular, interfacial, and solid-state systems, elemental diversity, diversity in both total charge and spin, molecular/structural diversity covering possible bonding and interaction motifs, and finally configurational diversity including both near-minima and far-from-minima configurations.

Regarding elemental diversity, most plane-wave DFT datasets follow the standards set by the Materials Project³³ and thus cover most of the periodic table up through plutonium. In contrast, molecular force field and DFT datasets are considerably more constrained in their elemental coverage, with most datasets only including C, H, N, and O and some coverage of S, F, Cl, and Br. The only molecular datasets of over 1 million data points with both energy and forces labels and which go beyond this set of elements are SPICE⁷⁸ (P, Li, Na, Mg, K, Ca, and I), AIMNet2 (P, I, Si, B, As, and Se), and more recently QCML (73 elements total)⁷⁹ and OMol25 (83 elements).⁸⁰ The only large-scale datasets with variable-charge species are SPICE(2),⁷⁸ AIMNet2,⁸¹ AIMNet-NSE,⁸² solvated protein fragments⁸³/GEMS,⁸⁴ and most recently QCML⁷⁹ and OMol2025.⁸⁰ While AIMNet-NSE⁸² is a dataset with variable spin species, it carries no force labels, whereas QCML does include both labels for some limited spin species⁷⁹ whereas variable spin is much better covered by OMol2025.⁸⁰ Another consideration for chemical diversity is the variety of chemical interactions captured by the data. While most of the molecular datasets mentioned thus far focus on isolated molecules, SPICE⁷⁸ makes an effort to capture varied intermolecular interactions between molecules, including some metal ions, solvated protein fragments⁸³/GEMS⁸⁴ includes protein fragments with multiple nearby solvent species, and ANI-2x⁸⁵ includes dimer scans and water clusters. Recognizing the value of intermolecular interactions, OMol25 also includes these data types.⁸⁰

In order to create structural diversity, large-scale data-acquisition efforts use a collection of sampling techniques such as normal modes, ab initio and empirical molecular dynamics, minimization and relaxation, reaction path sampling, metadynamics, and active learning, as outlined in Table 1. Structural diversity also encompasses the nature of the regions of the PES covered by the dataset. Having a broad range of forces between 0 eV/Å and ~ 10 eV/Å is important to effectively learn chemically-critical components of the PES encountered during MD, geometry optimization, or reaction path and transition state optimization. To that end, Transition-1x is the only large-scale dataset which explicitly samples reaction paths and transition states – critical areas for any truly foundational MLIP.^{86,87} Extremely high

Molecular Dataset	Size	Calculation method(s)	Sampling method(s)
OMol25 ⁸⁰	100M	ω B97M-D3/def2-TZVPD	Molecular dynamics Rattling Geometry optimization Reaction pathways
QCML ⁷⁹	33.5M	PBE0-D4/NAOs ⁸⁸	Chemical graphs Normal mode
AIMNet2 ⁸¹	20M	ω B97M-D3/def2-TZVPP	ANI-2x, ⁸⁵ OrbNet Denali ⁸⁹ Normal mode Metadynamics Molecular dynamics Torsion scan
∇^2 DFT ⁹⁰	16M	ω B97X-D/def2-SVP	Relaxation trajectory
Transition-1x ⁹¹	10M	ω B97X/6-31G*	Nudged elastic band
ANI-1x ⁹² /ANI-2x ⁸⁵	8.9M	ω B97X/6-31G*	Dimer and torsion scans Normal mode Molecular dynamics Active learning
QM7-X ⁹³	4.2M	PBE0-MBD	Normal mode
SPF ⁸³	2.7M	RPBE-D3(BJ)/def2-TZVP	Molecular dynamics Active learning
GEMS ⁸⁴	2.7M	PBE0-MBD/def2-TZVPP	SPF ⁸³ Molecular dynamics
SPICE ⁷⁸ /SPICE2 ⁹⁴	2M	ω B97M-D3(BJ)/def2-TZVPPD	Molecular dynamics Dimer scan
Materials Dataset	Size	Calculation method(s)	Sampling method(s)
OC20 ⁹⁵	265M	RPBE/PAW	Relaxation trajectory Molecular dynamics Rattling
OMat24 ⁶¹	110M	PBE(+U)/PAW	Rattled Boltzmann Molecular dynamics Relaxation trajectory
ODAC23 ⁹⁶	38M	PBE-D3(BJ)/PAW	Relaxation trajectory
Alexandria ⁹⁷	30M	PBE(+U)/PAW, SCAN/PAW PBEsol/PAW	Relaxation trajectory
OC22 ⁹⁸	9.8M	PBE(+U)/PAW	Relaxation trajectory
MPtrj ³⁴	1.6M	PBE(+U)/PAW	Relaxation trajectory
Unsupervised Dataset	Size	Calculation method(s)	Sampling method(s)
Zinc20 ⁷⁶ /Zinc22 ⁷⁷	4.5B	MMFF94	Torsion sampling
Uni-Mol ⁹⁹ /Uni-Mol2 ⁷⁵	838M	MMFF94	Distance geometry Geometry optimization

Table 1: Publicly available large scale datasets. We listed datasets of more than 1 million 3D structures with energy and force labels, as well as datasets of more than 100 million 3D structures without these labels.

forces (>50 eV/Å) can be detrimental to training and are usually filtered out, although high-energy and force data can be weighted with a Boltzmann factor,¹⁶ which still provides the model important chemical information about these regions. Finally, given that LLM FMs are pre-trained in an unsupervised manner on unlabeled data, it is worth asking what sort of data would be required to realize an analogous procedure in the chemical context. Ji and co-workers recently demonstrated that self-supervised pre-training on 800 million unlabeled 3D conformations of organic molecules improved down-stream performance on QM9 property prediction.⁷⁵




Figure 2: Proposed areas where substantial additional high-quality DFT data should be generated in order to realize a general and transferable foundational MLIP. a) Beyond-organic elements. Reproduced with permission from Ref.¹⁰⁰ b) Non-covalent interactions. Reproduced with permission from Ref.¹⁰¹ c) variable charge/spin/excited states. Adapted with permission from Ref.^{102,103} d) Complex reactivity. Reproduced with permission from Ref.¹⁰⁴ e) long-range interactions. Reproduced with permission from Ref.¹⁰⁵ f) Solid-solid interfaces. Reproduced with permission from Ref.¹⁰⁶ g) Surfaces under applied potential. Reproduced with permission from Ref.¹⁰⁷ h) Defects. Reproduced with permission from Ref.¹⁰⁸ i) Solid-liquid interfaces. Reproduced with permission from Ref.¹⁰⁹ j) Rich quantum information. Left: Reproduced with permission from Ref.¹¹⁰ Right: Reproduced with permission from Ref.¹¹¹

The specific areas where we believe the community should devote major effort to large-scale data generation are shown in Figure 2. For molecular DFT, the majority of the periodic table remains poorly covered, and under-representation of metal-organic complexes, complex

electrolytes, metalloenzymes, and disordered proteins underscores the need for substantial dataset generation efforts beyond organic elements (Fig. 2a). These types of systems could also address three other major gaps: non-covalent interactions (Fig. 2b), systems where long-range interactions are critical (Fig. 2g), and variable charge, spin, and excited states^{102,103} (Fig. 2c) which have been poorly covered by current datasets. Excited state energies and properties also offer an avenue for dataset development, where time-dependent DFT (TDDFT) calculations¹¹² are relatively inexpensive (roughly the cost of ground state DFT per state) and have been extensively benchmarked¹⁰² against higher level theories¹¹³ (Fig. 2e). We suggest that, beyond open-shell metals, combustion chemistry and atmospheric chemistry are spaces where different spin states are important and could be valuable. The final major gap we will describe is reactivity, particularly involving complex mechanisms and reactions in the condensed phase (Fig. 2d).

In the context of plane-wave DFT, we believe that major gaps include solid-solid interfaces (Fig. 2h), solid-liquid interfaces (Fig. 2j), surfaces under different applied potentials, (Fig. 2g), and defects (Fig. 2i). We note that some of these areas will require the use of better density functionals beyond the ubiquitous PBE and RPBE. Analogous to spin states for molecular properties, additional data generation for magnetic properties could also aid in correct magnetic moment initialization for diverse materials.

Finally, we believe there is a major need for large-scale datasets of rich quantum information, such as electron densities, orbital information, QTAIM descriptors, NBO interactions, and beyond (Fig. 2k). We note that infusing rich quantum information into ML models has recently been shown to be beneficial,^{110,111,114,115} and believe that powerful FM capabilities could be unlocked with such data. For example, MLIP FMs could be pre-trained on this extra information by swapping prediction schemes to predict atom and bond-level values. This could pave the way for foundational chemical models that go beyond MLIPs, serving as powerful representation models that can be adapted to predict forces and energies downstream.

The primary point of a pre-trained MLIP FMs is that they can be fine-tuned to yield accurate PESs for diverse chemical systems, or to enable property predictions beyond just energies and forces. At present, DFT calculations provide the best trade-off between accuracy and computational cost, such that nearly all supervised large-scale datasets for MLIPs are comprised of DFT calculations. There are at least three categories of errors in the reference values of such datasets, any one of which alone could justify fine-tuning. First is the choice of functional (or wavefunction "al" if one goes beyond DFT). Standard GGAs yield RMS errors that are typically 2-4 times larger than leading hybrid functionals^{2,116} which in turn have RMS errors that are at least 10 times larger than "gold standard" coupled cluster theory¹¹⁷ through perturbative triples, CCSD(T),¹¹⁸ or even beyond.¹¹⁹ This is a significant source of error. Turning from molecules to solids, there is an even greater need for high-quality data for fine-tuning. This begins with higher rung meta-GGA data in materials systems.¹²⁰⁻¹²³ The second error source is the choice of basis set; at least triple and preferably quadruple-zeta is required to converge hybrid DFT molecular relative energies towards the complete basis set (CBS) limit² yet many of the large-scale DFT datasets summarized in Table 1 use only double-zeta basis sets (e.g. 6-31G*) that induce significant errors. In condensed matter, the corresponding choice of AO basis or plane-wave cutoff is equally important.^{124,125}

Going beyond DFT is desirable. Here it is appropriate to highlight the effort of Smith and co-workers in generating the ANI-1ccx data set containing energies and forces of 500,000

organic molecules obtained with an accurate CCSD(T)*/CBS composite extrapolation.^{92,120} This protocol has deviations from the Schrödinger limit that are larger than normally acceptable for small, very high accuracy datasets,¹²⁶ but are much smaller than the best DFT errors, even in large basis sets.² Continued diverse collections of gold-standard CCSD(T) data (or beyond when necessary) would be extremely valuable for fine-tuning. Another result worth highlighting is the CCSD(T) assessment¹²⁷ of over 10,000 barrier heights in RDB7, which led to unexpectedly large RMS errors from the normally reliable ω B97X-D3 functional. The largest of these deviations were resolved by removing orbital instabilities and allowing the DFT to be spin-polarized when necessary.¹²⁸ However this highlights the limitations of even hybrid DFT for reactive systems where strong correlation effects are at play, as revealed by such symmetry-breaking.¹²⁹ For the most part, CC benchmarks for condensed matter periodic systems lie in the future, although there has been notable progress in CC methods for materials science.¹³⁰

The third factor are other error sources that prevent a given "model/basis" combination from yielding a well-defined value. In particular, DFT calculations on materials require careful attention to best practices^{125,131} to ensure reproducibility, beginning with choice of cutoffs, quadrature quality, and convergence criteria. Such factors have been discussed and investigated for the OC20 data.¹³² There are also non-trivial additional sources of error. Reaching the thermodynamic limit is particularly challenging for hybrid functionals.¹³³ Pseudopotential errors, typically not present in molecular calculations involving light elements, also require careful evaluation.¹³⁴⁻¹³⁶ Thresholds and cutoffs also affect molecular calculations, but can be controlled, and combined with SCF algorithms that ensure descent.¹³⁷ The local correlation methods needed for CCSD(T) on large molecules¹³⁸ are also subject to domain errors.¹³⁹ They scale with size, and may impact accuracy^{140,141} and so ideally should be removed and replaced by strict numerical thresholds.^{142,143} These systematic errors make dataset fusion difficult as machine learning models trained on such datasets will likely have difficulty learning physical properties/labels from disjoint underlying distributions. In addition, reviewers could help find inconsistencies before initial publication and scientists remain committed to maintaining datasets after publication. One recent example underscoring these realities was a study by Garrison et. al.¹⁴⁴ where the authors trained machine learning models on the tmQM¹⁴⁵ dataset but found many structures in the dataset to contain missing hydrogen atoms. After filtering and recomputing the dataset, model performance improved significantly.

Looking beyond energies and forces, accurate spectral simulations and other response properties which intimately depend on 3D atomic positions would be extremely valuable for use in fine-tuning, dramatically expanding the properties that fine-tuned FMs could predict. Finally, fine-tuning atomistic FMs with experimental data such as radial distribution functions, solvation free energies, density, and viscosity, could also be desirable. We note that experimental data can also be used in the context of selecting or "aligning" FMs trained on small cluster datasets to properties better reflected by bulk simulations. For example, Gong et al. map MLIP forces to pressures (and thus densities) from a small set of representative, experimental measurements.¹⁴⁶ We discuss additional training strategies that incorporate experimental data in Section 3.3.

3.3 Model-Agnostic Training Strategies

Effective training strategies have and will continue to play an important role in developing the next generation of FMs for chemistry and materials. We highlight multiple directions relevant to both pre-training as well as post-training schemes, with the goal of obtaining “model-agnostic” training strategies that can be applied to any foundational MLIP.

Pre-training strategies. While unsupervised training methods have been critical in improving the foundational capacity of LLMs and LVMs, in computational chemistry and materials science there are typically energy and/or forces labels (or molecular property labels for Zinc22⁷⁷/Zinc20⁷⁶) associated with atomistic configurations such that supervised learning is the common approach for MLIP optimization. By contrast self-supervised approaches start with the key ingredient being the availability of a data distribution that might sensibly center on molecular geometry, for instance, the PCQM4Mv2¹⁴⁷ and PM6¹⁴⁸ datasets developed under the PubChemQC Project.¹⁴⁹ This type of data is highly suitable for graph self-supervised learning using denoising that enables pre-training models to learn from molecular graphs for intrinsic chemical information. One recent example used geometric denoising to create a pre-trained model that exhibited fine-tuned accuracy improvement compared to randomly initialized "from scratch" models for various chemical properties in the QM9 dataset.^{55,150} This denoising approach was later extended to work for non-equilibrium structures.¹⁵¹ Understanding how much self-supervised pre-training relies on the quality of the training distribution, and how scalable such approaches can be in terms of data size compared to purely supervised (pre-)training, is currently an active area of investigation. An additional approach that is gaining acceptance is pre-training using synthetic energy and force data provided by other from-scratch trained MLIPs, which has been demonstrated to improve robustness of fine-tuned NN potentials on the condition that the synthetic source is reasonably reliable.¹⁵² Other examples of self-supervised learning approaches in materials chemistry include twin algorithms¹⁵³ and atom replacement¹⁵⁴ pre-training schemes.

Although most current models train on a dataset with calculations done at one level of theory,^{32,155} future foundational MLIPs could benefit from training on multiple levels of theory. A training objective that extracts learning signal from multiple levels of theory could take advantage of more abundant datasets generated with cheaper methods (like a semi-empirical potential), while still maintaining accuracy by leveraging higher fidelity data. Incorporating different levels of theory into the training pipeline would also allow MLIPs to learn about different interactions that are better described by different methods.^{31,156,157} Having multiple types of datasets also presents challenges in training a model that can handle different energy scales and the nuances of each level of theory. However, addressing these challenges could lead to the development of more general MLIPs.

Physical laws can also be incorporated into MLIPs by placing it directly into the loss function and training procedure, instead of making them inherent to the model architecture. For instance, rotational equivariance can be enforced via data augmentation loss as discussed previously, and energy conservation can be encouraged by minimizing the curl of the MLIP as part of the training objective.¹⁵⁸ While these “soft” constraints do not guarantee that the law will be obeyed, they circumvent the need to design new architectures for each new constraint of interest, and they can enable more scalable and flexible training.¹⁵⁹

Post-training strategies. After pre-training a FM, post-training is an important procedure

to further refine and optimize the model to a given downstream application. For domain specific tasks, such as energy and force prediction, modest amounts of higher-quality and labeled data are used within a standard supervised fine-tuning (SFT) context. Model distillation is another post-training approach for fine-tuning. In model distillation, a simpler model (the "student") is trained to replicate the behavior and performance of a larger, more complex model (the "teacher", or FM in this case). This process involves transferring knowledge from the teacher to the student by using the teacher’s predictions or internal representations as targets for training the student model. The use of teacher-student constructs frequently allows to perform more computationally efficient training, particularly for multi-label training cases.




Figure 3: Example of a post-training strategy for machine learning interatomic potentials via model distillation. One starts with an MLIP FM that has been trained on a large quantity of diverse data. A series of smaller MLIPs can then be trained via a knowledge distillation procedure, enabling specialization and speed for a specific task while retaining the general-purpose representations learned by the large-scale FM. Reproduced with permission from Ref.¹⁶⁰

Kelvinus et al. performed MLIP distillation by aligning node and edge features across models, but did not consider the task of specializing from general-purpose FMs.¹⁶¹ Recently Amin and co-workers¹⁶⁰ used distillation on large-scale models like MACE-OFF,¹⁵⁵ MACE-MP-0,³² and JMP.¹⁵⁶ By matching the Hessians of the energy predictions, these large-scale models were distilled into small, student MLIPs specialized for a specific chemical subset of the data and up to $50 \times$ faster during inference-time compared with the original large-scale model (see Fig. 3). This approach is agnostic to the MLIP architecture, and is compatible with arbitrary combinations of teacher and student MLIP design choices. Model distillation is one potential solution to balance scalable training with ensuring physical consistency in the final simulation. For instance,¹⁶⁰ showed that a large-scale model can be scalably trained without expensive constraints such as gradient-based forces for energy conservation, and can then be distilled into a student MLIP possessing gradient-based forces. This ensures

energy conservation while leveraging the general-purpose representations from the pre-trained large-scale model.

Thus far, much of the discussion has been about MLIPs trained on ab-initio data. However, experimental data is also a rich potential source of information to train MLIPs. Directly incorporating experimental data into MLIP training has primarily been done through differentiable simulation.^{71,162–166} In this case, MLIPs are typically initialized by pre-training with an analytical prior potential, or on a sparse set of simulated observables such as radial distribution functions, phonon densities, and diffusivity coefficients, which can be computed as ensemble averages over an equilibrium MD simulation. The pre-trained model is then fine-tuned to match experimental observables. Computing gradients necessary for MLIP finetuning through the full MD simulation is prohibitively expensive, and so techniques such as reweighting,¹⁶³ implicit differentiation,¹⁶⁷ and the adjoint method¹⁶⁸ are leveraged to enable tractable optimization. However, major technical challenges remain, the most notable of which is effectively training with *dynamical* observables.¹⁶⁹

3.4 Criteria for a Successful Foundation Model

A MLIP FM for atomistic simulation should perform well on a diverse range of tasks after fine-tuning and post-training, spanning different levels of theory, different types of systems, static and dynamic quantities, and a range of properties beyond just energies and forces. Unrelated downstream tasks might include field-dependent property predictions, generative models for crafting new structures or molecules,^{62,150,152,153,170} measuring energies and forces in defective structures,^{171,172} and dynamic testing environments created through MD simulations under more varied conditions.^{16,172–174}

One critical goal is the development of testing platforms for MLIP FMs to test their endowed capabilities. It demands considerable effort with regard to the computational cost of many simulation runs, studies to validate testing components, and analyses to derive scientific insights. But we believe that the benefits to the scientific community are profound and represent an essential direction for future research in demonstrating the promise of MLIP FMs.^{173,175–177} While competition over machine learning model performance on small datasets like MD17¹⁵⁸ did drive crucial early MLIP development, a number of large community efforts have emerged more recently that define a successive series of "leaderboards" as testing grounds for MLIPs. The OC20 leaderboard⁹⁵ was established in 2021, with a dataset split into training/validation/testing categories with rigorously defined in-distribution (ID) and out-of-distribution (OOD) systems, and a private test set to prevent test performance from being gamed. Energy and force model performance among methods developers on the OC20 leaderboard has consistently improved over time since its inception. In 2023, the MatBench Discovery leaderboard⁵⁸ was created to test MLIP performance for crystalline materials relaxations and the ability to predict materials thermodynamic stability and formation energy. While initially focused on the use of only MPtrj³⁴ as the training dataset, 2024 saw an explosion of interest in MatBench Discovery, including teams from Google, Microsoft, and Meta, each reporting the generation of supplemental training data as well as models trained on their data beating the previous best performing model. However, some of the new data may risk leakage with the test set. Further, most of the leaderboard evaluations are overly focused on energy, which may be overfit in order to improve leaderboard ranking while

degrading actual scientific utility.

A related issue is the lack of robust uncertainty quantification. The prevalent method for quantifying uncertainty in MLIP predictions involves using ensembles;^{178–180} this approach entails running multiple fits of the MLIP simultaneously and interpreting the variances among them as a measure of uncertainty. However, ensembles for quantifying uncertainty is computationally demanding, requiring parallelization, and may be unsuitable for MLIP FMs of larger scale. Such prediction uncertainty has not thus far been taken into account in previous leaderboards. Both issues have severe consequences in practical applications, and proper tests beyond simple accuracy measures have to be carefully designed.

Even if the MLIP universality within a specific domain can be evaluated through OOD tests, the leaderboards do not seek to show that the resulting models are foundational, i.e. that they obey scaling laws or are derived from diverse data or pre-training strategies or that they exhibit superior performance on a broad range of down-stream tasks after fine-tuning. While the large size of the OC20 training set makes it a viable arena in which to evaluate model scalability, e.g. via performance with 1M vs 10M vs 100M training data, only performance on the full training dataset is featured on the leaderboard. Further, both OC20 and MatBench Discovery focus on well-behaved relaxations but without consideration of other emergent properties such as MD stability or PES smoothness, which has been shown to be independent of just test errors for energies and forces.^{170,173,181,182} However the recently released NNP Arena¹⁸³ is a step in the right direction, by evaluating energy error for both pre-trained MLIPs (and pure DFT models) against gold-standard coupled cluster reference values, without any concern for what data was used in MLIP training or how MLIP training was performed. NNP Arena further evaluates inference speed, which previous leaderboards haven’t considered but which may be a very important consideration for the future. Another important direction pursued in the recent MLIP Arena¹⁸⁴ involves assessing the extent to which models comply with fundamental physics, including energy conservation, asymptotic behavior, smoothness, symmetry, and simulation stability. These metrics do not require additional labels, yet they carry substantial implications for the reliability and applicability of the models in practical settings.

3.5 Toward Foundation Models in Chemistry and Materials

At present all MLIPs are not true FMs by all criteria outlined in Section 2, but there are a few prototype MLIP FMs on the horizon that are proving more accurate, to more comprehensively cover the periodic table for molecular systems, and also show the ability to comply with heuristic scaling laws. In particular the pretraining models such as Efficiently Scaled Attention Interatomic Potential (EScAIP)⁴⁴ and Uni-Mol2⁷⁵ have characterized the scaling correlations of decreasing validation error with increasing model size, dataset size, and computational resources. This would suggest that emergent capabilities might follow, such as predicting energy labels from unlabeled data as shown for Uni-Mol2.⁷⁵ The fact that LLMs have demonstrated such scaling laws has led to emergent properties in chemistry for generative tasks in small molecule drugs⁷ and their synthesis,⁹ although the LLMs were not trained on SMILES strings or chemical reaction templates.

Since the initial writing of this manuscript, two important contributions have been released which substantially advance the state-of-the-art for atomistic simulation: the Open Molecules

2025 (OMol25) dataset,⁸⁰ and the Universal Model of Atoms pre-trained MLIP.¹⁸⁵ OMol25 is the largest ever molecular DFT dataset for training MLIPs, with over 100 million snapshots at an excellent level of DFT (ω B97M-V/def2-TZVPD), where snapshots include up to 350 atoms, up to ten unpaired electrons, total charges ranging from -10e to +10e, and 83 elements of the periodic table. OMol25 broadly covers small molecules, electrolytes, biomolecules, and metal complexes, using more than ten distinct methods for structural sampling.⁸⁰ OMol25 thus makes substantial progress towards high-quality data coverage beyond organics, such as non-covalent interactions, spin and charge species, complex reactions, and long-range interactions, nearly half of the areas outlined in Figure 2. In addition to the >100 million training data, OMol25 also includes explicit out-of-distribution test sets and novel model evaluations, which will populate a public leaderboard in the near future, and baseline trained MLIPs.

Simultaneous with OMol25’s release was the release of UMA,¹⁸⁵ which is trained on OMol25,⁸⁰ OMat24,⁶¹ ODAC23,⁹⁶ OC20,¹³² and the novel Open Molecular Crystal dataset (unpublished), totaling nearly 500 million training data - the largest and most diverse training set ever used to train a single MLIP. UMA also demonstrates scaling laws for both number of learnable parameters and compute, thus coming the closest to our vision of a foundation model for atomistic simulation. UMA is based on the eSEN architecture¹⁸⁶ and further employs a novel mixture of linear experts (MoLE) for two key benefits: 1) the MoLE causes a model trained on all datasets simultaneously, in a multi-task fashion (where each distinct training dataset is a distinct "task"), to exhibit superior performance to models trained separately on each dataset. 2) the MoLE allows a huge number of learnable parameters to be active during training, but a much smaller number to be active at inference time, allowing the model to effectively leverage the vast training dataset while still remaining fairly fast and memory efficient for subsequent simulations.

UMA and eSEN-OMol now top Rowan’s NNP Arena leaderboard¹⁸³ for neutral, closed-shell, organic molecular energies, with errors against coupled cluster for the GMTKN55 benchmark set falling by nearly 5 kcal/mol compared with the previous best model. Digging more deeply into the GMTKN55 test set,¹⁸⁷ one can see that UMA’s energies are already lower error with respect to CCSD(T) than ω B97M-D3BJ/def2-QZVP DFT for variable charge, variable spin, and metal-containing large systems and barrier heights, while being nearly comparable for intermolecular non-covalent interactions. However, the OMol25 evaluation tasks demonstrate that there is still substantial room for improvement on charge, spin, and long-range interactions, where error against reference DFT data remains very large (> 100 meV) likely due to the naive handling of charge and spin and the 6Å graph cutoff, emphasizing the key areas where further MLIP architecture development is needed.

4 Discussion and Future Directions

The primary promise of a MLIP foundation model is that it would enable massively impactful investigation of systems ranging from metal-organic frameworks to enzymes to electrocatalysis beyond current physics-only models, while avoiding the need to build each ML model from the ground-up for any new chemical system of interest. One key insight that has been observed in other areas of machine learning is that methods that may work in “low data”

regimes do not necessarily work as well in higher data regimes. As data size increases, as we have summarized for the many large-scale community and data efforts in chemistry and materials, can ML architectures scale accordingly to handle such data? Currently, the two paradigms are to use inductive biases to address physical constraints, often requiring less data,^{41,42} with an alternative philosophy being that such physical laws can be easily learned when there is enough data, and when model capacity and compute is correspondingly scaled to take advantage of this data. Constrained MLIPs or learning energy and forces together ensures that conservative forces are available for the primary downstream MD task,¹⁵⁹ but post-training strategies can take non-conservative MLIPs and distill them into conservative MLIPs,¹⁶⁰ enabling both initial training at scale and physical soundness at deployment. Hence we posit here that incorporation of architectural inductive biases is not the only way to create sound MLIP FMs, a direction that is currently an underexplored paradigm in chemistry and materials. It raises the interesting question whether the chemical and materials sciences will escape the "bitter lesson", or whether our science domains may also conclude that data and scalability is superior to built-in constraints, as has come to pass in other fields such as robotics and weather that are also grounded in physics.

Training a truly foundational MLIP that can realize the far-reaching impact we envision will require a collection of very large datasets that broadly span chemistry, materials, phases, and interfaces. While our earlier discussion focused on using DFT-labeled data as the standard approach for pre-training, perhaps even much larger low-level datasets with or without labels from semi-empirical methods or classical force fields or on-the-fly schemes for computing properties could be important to unlock powerful at-scale capabilities for chemical FMs. It is also possible that the best approach will not actually start with an MLIP, but will instead leverage some combination of traditional chemoinformatics, chemically-relevant text, and/or un-labeled structures to create datasets suited for larger FMs which can be retooled as MLIPs. Information contained in other multi-modal data that incorporates experimental information could also be extremely valuable for pre-training, although realizing this in practice will require massive data curation efforts in order for the scale of the data to be suitable for pre-training as well as the development of novel training strategies even beyond what we have discussed to be viable at scale.

While current MLIP testing platforms do provide a substantial resource to the community, there are still significant challenges and opportunities in the context of MLIP FMs. For example, there is not always a direct correlation between reduced force errors observed in test set performance and its actual performance in MD simulations.^{170,173,181,182} Current MLIP leaderboards also do not yet explicitly evaluate model scalability with respect to the amount of training data or learnable parameters, nor do they seek to quantify the "foundational capacity" of the model via standardized fine-tuning across chemical and materials domains and properties. Scaling laws being the defining property of a FM, MLIPs should demonstrate continuous improvement with more parameters, more data, and more compute. Ultimately the FM practitioner must demonstrate the success of a FM via fine-tuned models that are fast and specialized - and which are more accurate than a model trained from scratch on data specific to that task/domain. The development of a scalable FM will also require using efficient algorithms to perform pre-training and fine-tuning. In particular, efficient mechanism for storing and retrieving data for batched training, dynamic load balancing methods that can accommodate variations in graph sizes, and communication hiding techniques that overlap

computation and message passing will be important for achieving a rapid turnaround.

We note that the success of foundational LLMs and LVMs has critically depended on acquisition of vast collections of text and images from the internet, i.e. data obtained via implied or enforced open access. But in the context of DFT data for training MLIPs, this assumption of open access to data is not always adhered to when commercial players are involved. For instance, the datasets used to train Google’s GNoME model for crystal energies and Microsoft Research reported MatterSim have not been released for public use. In contrast, the FAIR Chemistry team at Meta, responsible for the largest MLIP training datasets built to-date, has consistently released their datasets (and trained models) fully open-source, a value to the scientific community that cannot be understated. In a related vein, MLIP models have benefited tremendously from industry-released software tools such as TensorFlow,¹⁸⁸ Keras,¹⁸⁹ and PyTorch.¹⁹⁰ However, open-access development has not occurred for LLMs such as ChatGPT,¹⁹¹ and the loss of implementation details on how LLM FMs work could limit our ability to develop similarly powerful MLIP FMs. We note that while DeepSeek-V3 is open source, and can deliver high-quality results at a lower cost compared to close-source LLMs, it is either restricted or outright banned in several countries due to concerns about ethics, privacy, and security. As is needed for all emerging technologies, training in ethical development and use of AI/ML should remain a top priority in the physical sciences.¹⁹² An additional practical challenge is the computational resources required for the demonstration of FM capabilities at large scale. It will require world-class high performance computing beyond the academic lab scale and the first place to turn to are the leadership computing levels at the national labs. The combined issues of open source data, software, models, politics, and ethics training will likely require an even larger federation of government agencies with industry, and international cooperation, modeled on the paradigm of scale of the Human Genome Project.¹⁹³

References

- ¹ Pople, J. A. *Theoretical models for chemistry*, 51–61 (John Wiley and Sons, 1973). URL <https://inis.iaea.org/records/jrnfw-ccm20>.
- ² Mardirossian, N. & Head-Gordon, M. Thirty years of density functional theory in computational chemistry: an overview and extensive assessment of 200 density functionals. *Mol. Phys.* **115**, 2315–2372 (2017). URL <https://www.tandfonline.com/doi/abs/10.1080/00268976.2017.1333644>.
- ³ Hestness, J. *et al.* Deep learning scaling is predictable, empirically. *arXiv preprint arXiv:1712.00409* (2017). URL <https://arxiv.org/abs/1712.00409v1>.
- ⁴ Bommasani, R. *et al.* On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258* (2021). URL <https://crfm.stanford.edu/assets/report.pdf>.
- ⁵ Zheng, Z. *et al.* Chatgpt research group for optimizing the crystallinity of mofs and cofs. *ACS Cent. Sci.* **9**, 2161–2170 (2023). URL <https://pubs.acs.org/doi/full/10.1021/acscentsci.3c01087>.

- ⁶ Boiko, D. A., MacKnight, R., Kline, B. & Gomes, G. Autonomous chemical research with large language models. *Nature* **2023** *624*:7992 **624**, 570–578 (2023). URL <https://www.nature.com/articles/s41586-023-06792-0>.
- ⁷ Cavanagh, J. M. *et al.* Smileyllama: Modifying large language models for directed chemical space exploration. *arXiv preprint arXiv:2409.02231* (2024). URL <https://arxiv.org/abs/2409.02231v2>.
- ⁸ Bran, A. M. *et al.* Augmenting large language models with chemistry tools. *Nat. Mach. Intell.* **2024** *6*:5 **6**, 525–535 (2024). URL <https://www.nature.com/articles/s42256-024-00832-8>.
- ⁹ Sun, K. *et al.* Synllama: Generating synthesizable molecules and their analogs with large language models. *Arxiv* **2503.12602** (2025).
- ¹⁰ Unke, O. T. *et al.* Machine Learning Force Fields. *Chem. Rev.* **121**, 10142–10186 (2021). URL <https://pubs.acs.org/doi/10.1021/acs.chemrev.0c01111>.
- ¹¹ Wang, G. *et al.* Machine learning interatomic potential: Bridge the gap between small-scale models and realistic device-scale simulations. *iScience* **27**, 109673 (2024). URL <https://doi.org/10.1016/j.isci.2024.109673><https://linkinghub.elsevier.com/retrieve/pii/S2589004224008952>.
- ¹² Allen, A. E. A. *et al.* Learning together: Towards foundation models for machine learning interatomic potentials with meta-learning. *npj Comput. Mater.* **10**, 154 (2024). URL <http://dx.doi.org/10.1038/s41524-024-01339-x><https://www.nature.com/articles/s41524-024-01339-x>.
- ¹³ Smith, J. S., Nebgen, B., Lubbers, N., Isayev, O. & Roitberg, A. E. Less is more: Sampling chemical space with active learning. *The Journal of Chemical Physics* **148**, 241733 (2018). URL <https://doi.org/10.1063/1.5023802>.
- ¹⁴ Khalak, Y., Tresadern, G., Hahn, D. F., de Groot, B. L. & Gapsys, V. Chemical space exploration with active learning and alchemical free energies. *Journal of Chemical Theory and Computation* **18**, 6259–6270 (2022). URL <https://doi.org/10.1021/acs.jctc.2c00752>.
- ¹⁵ Kulichenko, M. *et al.* Uncertainty-driven dynamics for active learning of interatomic potentials. *Nature Computational Science* **3**, 230–239 (2023).
- ¹⁶ Guan, X., Heindel, J. P., Ko, T., Yang, C. & Head-Gordon, T. Using machine learning to go beyond potential energy surface benchmarking for chemical reactivity. *Nat. Comp. Sci.* **2023** *3*:11 **3**, 965–974 (2023). URL <https://www.nature.com/articles/s43588-023-00549-5>.
- ¹⁷ Nandy, A. From pages to patterns: Towards extracting catalytic knowledge from structure and text for transition-metal complexes and metal-organic frameworks. *J. Catal.* (2025).

- ¹⁸ Yao, L., Ou, Z., Luo, B., Xu, C. & Chen, Q. Machine learning to reveal nanoparticle dynamics from liquid-phase tem videos. *ACS Central Science* **6**, 1421–1430 (2020). URL <https://doi.org/10.1021/acscentsci.0c00430>.
- ¹⁹ Brown, T. *et al.* Language models are few-shot learners. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M. & Lin, H. (eds.) *Advances in Neural Information Processing Systems*, vol. 33, 1877–1901 (Curran Associates, Inc., 2020). URL https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfcb4967418bfb8ac142f64a-Paper.pdf.
- ²⁰ Touvron, H. *et al.* Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971* (2023). URL <https://arxiv.org/abs/2302.13971v1>.
- ²¹ Kaplan, J. *et al.* Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361* (2020). URL <https://arxiv.org/abs/2001.08361v1>.
- ²² Hoffmann, J. *et al.* Training compute-optimal large language models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems, NIPS '22* (Curran Associates Inc., Red Hook, NY, USA, 2022).
- ²³ Bahri, Y., Dyer, E., Kaplan, J., Lee, J. & Sharma, U. Explaining neural scaling laws. *Proc. Natl. Acad. Sci. USA* **121**, e2311878121 (2024). URL <https://doi.org/10.1073/pnas.2311878121>.
- ²⁴ Vaswani, A. *et al.* Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17*, 6000–6010 (Curran Associates Inc., Red Hook, NY, USA, 2017).
- ²⁵ Brohan, A. *et al.* RT-1: Robotics Transformer for Real-World Control at Scale. In *Proceedings of Robotics: Science and Systems* (Daegu, Republic of Korea, 2023).
- ²⁶ Zitkovich, B. *et al.* RT-2: Vision-language-action models transfer web knowledge to robotic control. In *7th Annual Conference on Robot Learning* (2023). URL <https://openreview.net/forum?id=XMQgwiJ7KSX>.
- ²⁷ Kim, M. J. *et al.* Openvla: An open-source vision-language-action model. In Agrawal, P., Kroemer, O. & Burgard, W. (eds.) *Proceedings of The 8th Conference on Robot Learning*, vol. 270 of *Proc. Mach. Learn. Res.*, 2679–2713 (PMLR, 2025). URL <https://proceedings.mlr.press/v270/kim25c.html>.
- ²⁸ Ghosh, D. *et al.* Octo: An Open-Source Generalist Robot Policy. In *Proceedings of Robotics: Science and Systems* (Delft, Netherlands, 2024).
- ²⁹ Wei, J. *et al.* Emergent abilities of large language models. *Transactions on Machine Learning Research (TMLR)* (2022). URL <https://arxiv.org/abs/2206.07682v2>.
- ³⁰ Schaeffer, R., Miranda, B. & Koyejo, S. Are emergent abilities of large language models a mirage? *Advances in Neural Information Processing Systems (NeurIPS)* (2023). URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/ad98a266f45005c403b8311ca7e8bd7-Abstract-Conference.html.

- ³¹ Zhang, D. *et al.* DPA-2: a large atomic model as a multi-task learner. *npj Comput. Mater.* **2024** *10:1* **10**, 1–15 (2024). URL <https://www.nature.com/articles/s41524-024-01493-2>.
- ³² Batatia, I. *et al.* A foundation model for atomistic materials chemistry. *arXiv preprint arXiv:2401.00096* (2023). URL <https://arxiv.org/abs/2401.00096v2>.
- ³³ Jain, A. *et al.* Commentary: The materials project: A materials genome approach to accelerating materials innovation. *APL Mater.* **1**, 11002 (2013). URL [/aip/apm/article/1/1/011002/119685/Commentary-The-Materials-Project-A-materials](http://aip/apm/article/1/1/011002/119685/Commentary-The-Materials-Project-A-materials).
- ³⁴ Deng, B. *et al.* CHGNet as a pretrained universal neural network potential for charge-informed atomistic modelling. *Nat. Mach. Intell.* **2023** *5:9* **5**, 1031–1041 (2023). URL <https://www.nature.com/articles/s42256-023-00716-3>.
- ³⁵ Kaur, H. *et al.* Data-efficient fine-tuning of foundational models for first-principles quality sublimation enthalpies. *Faraday Discussions* **256**, 120–138 (2025). URL <https://pubs.rsc.org/en/content/articlehtml/2025/fd/d4fd00107ahttps://pubs.rsc.org/en/content/articlelanding/2025/fd/d4fd00107a>.
- ³⁶ Behler, J. & Parrinello, M. Generalized neural-network representation of high-dimensional potential-energy surfaces. *Phys. Rev. Lett.* **98**, 146401 (2007). URL <https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.98.146401>.
- ³⁷ Zhang, L., Han, J., Wang, H., Car, R. & Weinan, E. Deep potential molecular dynamics: A scalable model with the accuracy of quantum mechanics. *Phys. Rev. Lett.* **120**, 143001 (2018). URL <https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.120.143001>.
- ³⁸ Drautz, R. Atomic cluster expansion for accurate and transferable interatomic potentials. *Phys. Rev. B* **99**, 014104 (2019). URL <https://journals.aps.org/prb/abstract/10.1103/PhysRevB.99.014104>.
- ³⁹ Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O. & Dahl, G. E. Neural message passing for quantum chemistry. *34th International Conference on Machine Learning, ICML 2017* **3**, 2053–2070 (2017). URL <https://arxiv.org/abs/1704.01212v2>.
- ⁴⁰ Lubbers, N., Smith, J. S. & Barros, K. Hierarchical modeling of molecular energies using a deep neural network. *The Journal of Chemical Physics* **148**, 241715 (2018). URL <https://doi.org/10.1063/1.5011181>.
- ⁴¹ Batzner, S. *et al.* E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nat. Comm.* **2022** *13:1* **13**, 1–11 (2021). URL <http://arxiv.org/abs/2101.03164http://dx.doi.org/10.1038/s41467-022-29939-5>.
- ⁴² Haghighatlari, M. *et al.* Newtonnet: a newtonian message passing network for deep learning of interatomic potentials and forces. *Digit. Discov.* **1**, 333–343 (2022). URL <https://pubs.rsc.org/en/content/articlehtml/2022/dd/d2dd00008chttps://pubs.rsc.org/en/content/articlelanding/2022/dd/d2dd00008c>.

- ⁴³ Liao, Y. L. & Smidt, T. Equiformer: Equivariant graph attention transformer for 3d atomistic graphs. *11th International Conference on Learning Representations, ICLR 2023* (2022). URL <https://arxiv.org/abs/2206.11990v2>.
- ⁴⁴ Qu, E. & Krishnapriyan, A. S. The importance of being scalable: Improving the speed and accuracy of neural network interatomic potentials across chemical domains. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems* (2024). URL <https://openreview.net/forum?id=Y4mBaZu4vy>.
- ⁴⁵ Liu, S. *et al.* Multiresolution 3d-densenet for chemical shift prediction in nmr crystallography. *J. Phys. Chem. Lett.* **10**, 4558–4565 (2019). URL <https://doi.org/10.1021/acs.jpclett.9b01570>.
- ⁴⁶ Musil, F. *et al.* Physics-inspired structural representations for molecules and materials. *Chem. Rev.* **121**, 9759–9815 (2021). URL <https://pubs.acs.org/doi/full/10.1021/acs.chemrev.1c00021>.
- ⁴⁷ Sutton, R. The bitter lesson (2019). URL <http://www.incompleteideas.net/IncIdeas/BitterLesson.html>.
- ⁴⁸ Márquez-Neila, P., Salzmann, M. & Fua, P. Imposing hard constraints on deep networks: Promises and limitations. *arXiv preprint arXiv:1706.02025* (2017).
- ⁴⁹ Dosovitskiy, A. *et al.* An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations* (2021). URL <https://openreview.net/forum?id=YicbFdNTTy>.
- ⁵⁰ Gruver, N., Finzi, M. A., Goldblum, M. & Wilson, A. G. The lie derivative for measuring learned equivariance (2023). URL <https://openreview.net/forum?id=JL7Va5Vy15J>.
- ⁵¹ Grattafiori, A. *et al.* The llama 3 herd of models. *arXiv preprint arXiv:2407.21783* (2024). URL <https://arxiv.org/abs/2407.21783>.
- ⁵² Yu, X. *et al.* Point-bert: Pre-training 3d point cloud transformers with masked point modeling (2022). URL <https://arxiv.org/abs/2111.14819>. 2111.14819.
- ⁵³ Abramson, J. *et al.* Accurate structure prediction of biomolecular interactions with alphafold 3. *Nature* **630**, 493–500 (2024). URL <http://dx.doi.org/10.1038/s41586-024-07487-w>.
- ⁵⁴ Liao, Y. L., Wood, B., Das, A. & Smidt, T. Equiformerv2: Improved equivariant transformer for scaling to higher-degree representations. *12th International Conference on Learning Representations, ICLR 2024* (2023). URL <https://arxiv.org/abs/2306.12059v3>.
- ⁵⁵ Neumann, M. *et al.* Orb: A fast, scalable neural network potential. *arXiv preprint arXiv:2410.22570* (2024). URL <https://arxiv.org/abs/2410.22570v1>.

- ⁵⁶ Bigi, F., Langer, M. F. & Ceriotti, M. The dark side of the forces: assessing non-conservative force models for atomistic machine learning. *arXiv preprint arXiv:2412.11569* (2024). URL <https://arxiv.org/abs/2412.11569v1>.
- ⁵⁷ Wang, H.-C., Botti, S. & Marques, M. A. L. Predicting stable crystalline compounds using chemical similarity. *npj Comput. Mater.* **7**, 12 (2021).
- ⁵⁸ Riebesell, J. *et al.* Matbench discovery – a framework to evaluate machine learning crystal stability predictions. *arXiv preprint arXiv:2308.14920* (2023). URL <https://arxiv.org/abs/2308.14920v3>.
- ⁵⁹ Kreiman, T. & Krishnapriyan, A. S. Understanding and Mitigating Distribution Shifts For Machine Learning Force Fields. *arXiv preprint arXiv:2503.08674* (2025). URL <https://arxiv.org/abs/2503.08674>. ArXiv:2503.08674 [cs].
- ⁶⁰ Deng, B. *et al.* Systematic softening in universal machine learning interatomic potentials. *npj Comput. Mater.* **11**, 9 (2025). URL <http://arxiv.org/abs/2405.07105><https://www.nature.com/articles/s41524-024-01500-6>.
- ⁶¹ Barroso-Luque, L. *et al.* Open materials 2024 (omat24) inorganic materials dataset and models. *arXiv preprint arXiv:2410.12771* (2024). URL <https://arxiv.org/abs/2410.12771v1>.
- ⁶² Yang, H. *et al.* Mattersim: A deep learning atomistic model across elements, temperatures and pressures. *arXiv preprint arXiv:2405.04967* (2024). URL <https://arxiv.org/abs/2405.04967v2>.
- ⁶³ Mazitov, A. *et al.* Pet-mad, a universal interatomic potential for advanced materials modeling. *arXiv preprint arXiv:2503.14118* (2025).
- ⁶⁴ Yue, S. *et al.* When do short-range atomistic machine-learning models fall short? *J. Chem. Phys.* **154**, 34111 (2021). URL [/aip/jcp/article/154/3/034111/199865/When-do-short-range-atomistic-machine-learning](http://aip/jcp/article/154/3/034111/199865/When-do-short-range-atomistic-machine-learning).
- ⁶⁵ Niblett, S. P., Galib, M. & Limmer, D. T. Learning intermolecular forces at liquid-vapor interfaces. *J. Chem. Phys.* **155**, 164101 (2021). URL [/aip/jcp/article/155/16/164101/199859/Learning-intermolecular-forces-at-liquid-vapor](http://aip/jcp/article/155/16/164101/199859/Learning-intermolecular-forces-at-liquid-vapor).
- ⁶⁶ Rodgers, J. M. & Weeks, J. D. Interplay of local hydrogen-bonding and long-ranged dipolar forces in simulations of confined water. *Proc. Natl. Acad. Sci. USA* **105**, 19136–19141 (2008). URL <https://www.pnas.org/doi/abs/10.1073/pnas.0807623105>.
- ⁶⁷ Cox, S. J. Dielectric response with short-ranged electrostatics. *Proc. Natl. Acad. Sci. USA* **117**, 19746–19752 (2020). URL <https://www.pnas.org/doi/abs/10.1073/pnas.2005847117>.
- ⁶⁸ Grisafi, A. & Ceriotti, M. Incorporating long-range physics in atomic-scale machine learning. *J. Chem. Phys.* **151**, 204105 (2019). URL [/aip/jcp/article/151/20/204105/197247/Incorporating-long-range-physics-in-atomic-scale](http://aip/jcp/article/151/20/204105/197247/Incorporating-long-range-physics-in-atomic-scale).

- ⁶⁹ Cheng, B. Latent ewald summation for machine learning of long-range interactions. *arXiv preprint arXiv:2408.15165* (2024). URL <https://arxiv.org/abs/2408.15165v2>.
- ⁷⁰ Kim, D., King, D. S., Zhong, P. & Cheng, B. Learning charges and long-range interactions from energies and forces. *arXiv preprint arXiv:2412.15455* (2024).
- ⁷¹ Raja, S., Amin, I., Pedregosa, F. & Krishnapriyan, A. S. Stability-aware training of machine learning force fields with differentiable boltzmann estimators. *Transactions on Machine Learning Research* (2025). URL <https://openreview.net/forum?id=ZckLMG00s0>.
- ⁷² Giovanni, F. D. *et al.* On over-squashing in message passing neural networks: The impact of width, depth, and topology. *Proc. Mach. Learn. Res.* **202**, 7865–7885 (2023). URL <https://arxiv.org/abs/2302.02941v3>.
- ⁷³ Rusch, T. K., Bronstein, M. M. & Mishra, S. A survey on oversmoothing in graph neural networks. *arXiv preprint arXiv:2303.10993* (2023). URL <https://arxiv.org/abs/2303.10993v1>.
- ⁷⁴ Sriram, A., Das, A., Wood, B. M., Goyal, S. & Zitnick, C. L. Towards training billion parameter graph neural networks for atomic simulations. *ICLR 2022 - 10th International Conference on Learning Representations* (2022). URL <https://arxiv.org/abs/2203.09697v1>.
- ⁷⁵ Ji, X. *et al.* Uni-mol2: Exploring molecular pretraining model at scale. *Advances in Neural Information Processing Systems (NeurIPs)* (2024). URL <https://arxiv.org/abs/2406.14969v2>.
- ⁷⁶ Irwin, J. J. *et al.* Zinc20—a free ultralarge-scale chemical database for ligand discovery. *J. Chem. Info. Model.* **60**, 6065–6073 (2020). URL <https://doi.org/10.1021/acs.jcim.0c00675>.
- ⁷⁷ Tingle, B. I. *et al.* Zinc-22—a free multi-billion-scale database of tangible compounds for ligand discovery. *J. Chem. Info. Model.* **63**, 1166–1176 (2023). URL <https://doi.org/10.1021/acs.jcim.2c01253>.
- ⁷⁸ Eastman, P. *et al.* Spice, a dataset of drug-like molecules and peptides for training machine learning potentials. *Sci. Data 2022 10:1* **10**, 1–11 (2023). URL <https://www.nature.com/articles/s41597-022-01882-6>.
- ⁷⁹ Gansch, S. *et al.* The qcml dataset, quantum chemistry reference data from 33.5m dft and 14.7b semi-empirical calculations. *Scientific Data* **12**, 406 (2025). URL <https://doi.org/10.1038/s41597-025-04720-7>.
- ⁸⁰ Levine, D. S. *et al.* The open molecules 2025 (omol25) dataset, evaluations, and models (2025). URL <https://arxiv.org/abs/2505.08762>. 2505.08762.
- ⁸¹ Anstine, D. M., Zubatyuk, R. & Isayev, O. Aimnet2: A neural network potential to meet your neutral, charged, organic, and elemental-organic needs. *ChemRxiv* (2024). URL <https://chemrxiv.org/engage/chemrxiv/article-details/6763b51281d2151a022fb6a5>.

- ⁸² Zubatyuk, R., Smith, J. S., Nebgen, B. T., Tretiak, S. & Isayev, O. Teaching a neural network to attach and detach electrons from molecules. *Nat. Comm.* **2021 12:1** **12**, 1–11 (2021). URL <https://www.nature.com/articles/s41467-021-24904-0>.
- ⁸³ Unke, O. T. & Meuwly, M. Physnet: A neural network for predicting energies, forces, dipole moments, and partial charges. *J. Chem. Theo. Comput.* **15**, 3678–3693 (2019). URL <https://pubs.acs.org/doi/full/10.1021/acs.jctc.9b00181>.
- ⁸⁴ Unke, O. T. *et al.* Biomolecular dynamics with machine-learned quantum-mechanical force fields trained on diverse chemical fragments. *Sci. Adv.* **10** (2024). URL <https://www.science.org/doi/10.1126/sciadv.adn4397>.
- ⁸⁵ Devereux, C. *et al.* Extending the applicability of the ani deep learning molecular potential to sulfur and halogens. *J. Chem. Theo. Comput.* **16**, 4192–4202 (2020). URL <https://pubs.acs.org/doi/full/10.1021/acs.jctc.0c00121>.
- ⁸⁶ Yuan, E. C. *et al.* Analytical ab initio hessian from a deep learning potential for transition state optimization. *Nat. Comm.* **2024 15:1** **15**, 1–9 (2024). URL <https://www.nature.com/articles/s41467-024-52481-5>.
- ⁸⁷ Wander, B., Shuaibi, M., Kitchin, J. R., Ulissi, Z. W. & Zitnick, C. L. Cattsunami: Accelerating transition state energy calculations with pre-trained graph neural networks. *arXiv preprint arXiv:2405.02078* (2024). URL <https://arxiv.org/abs/2405.02078v3>.
- ⁸⁸ Blum, V. *et al.* Ab initio molecular simulations with numeric atom-centered orbitals. *Computer Physics Communications* **180**, 2175–2196 (2009). URL <https://www.sciencedirect.com/science/article/pii/S0010465509002033>.
- ⁸⁹ Christensen, A. S. *et al.* Orbnet denali: A machine learning potential for biological and organic chemistry with semi-empirical cost and dft accuracy. *J. Chem. Phys.* **155**, 204103 (2021). URL [/aip/jcp/article/155/20/204103/280547/OrbNet-Denali-A-machine-learning-potential-for](https://aip/jcp/article/155/20/204103/280547/OrbNet-Denali-A-machine-learning-potential-for).
- ⁹⁰ Khrabrov, K. *et al.* ∇^2 dft: A universal quantum chemistry dataset of drug-like molecules and a benchmark for neural network potentials. *Advances in Neural Information Processing Systems (NeurIPs)* (2024). URL <https://arxiv.org/abs/2406.14347v2>.
- ⁹¹ Schreiner, M., Bhowmik, A., Vegge, T., Busk, J. & Winther, O. Transition1x - a dataset for building generalizable reactive machine learning potentials. *Sci. Data* **2022 9:1** **9**, 1–9 (2022). URL <https://www.nature.com/articles/s41597-022-01870-w>.
- ⁹² Smith, J. S. *et al.* The ani-1ccx and ani-1x data sets, coupled-cluster and density functional theory properties for molecules. *Sci. Data* **2020 7:1** **7**, 1–10 (2020). URL <https://www.nature.com/articles/s41597-020-0473-z>.
- ⁹³ Hoja, J. *et al.* Qm7-x, a comprehensive dataset of quantum-mechanical properties spanning the chemical space of small organic molecules. *Sci. Data* **2021 8:1** **8**, 1–11 (2021). URL <https://www.nature.com/articles/s41597-021-00812-2>.

- ⁹⁴ Eastman, P., Pritchard, B. P., Chodera, J. D. & Markland, T. E. Nutmeg and spice: Models and data for biomolecular machine learning. *J. Chem. Theory Comput.* **20**, 8583–8593 (2024).
- ⁹⁵ Chanussot, L. *et al.* Open catalyst 2020 (oc20) dataset and community challenges. *ACS Catalysis* **11**, 6059–6072 (2021). URL <https://pubs.acs.org/doi/full/10.1021/acscatal.0c04525>.
- ⁹⁶ Sriram, A. *et al.* The open dac 2023 dataset and challenges for sorbent discovery in direct air capture. *ACS Cent. Sci.* **10**, 923–941 (2024). URL <https://pubs.acs.org/doi/full/10.1021/acscentsci.3c01629>.
- ⁹⁷ Schmidt, J. *et al.* Improving machine-learning models in materials science through large datasets. *Mat. Today Phys.* **48**, 101560 (2024).
- ⁹⁸ Tran, R. *et al.* The open catalyst 2022 (oc22) dataset and challenges for oxide electrocatalysts. *AICHE Annual Meeting, Conference Proceedings 2023-November*, 3066–3084 (2023). URL <https://pubs.acs.org/doi/full/10.1021/acscatal.2c05426>.
- ⁹⁹ Zhou, G. *et al.* Uni-mol: A universal 3d molecular representation learning framework (2018). URL <http://arxiv.org/abs/1802.08219>.
- ¹⁰⁰ Holzer, C., Gordiy, I., Grimme, S. & Bursch, M. Hybrid dft geometries and properties for 17k lanthanoid complexes-the lmqm data set. *Journal of Chemical Information and Modeling* **64**, 825–836 (2024). URL <http://dx.doi.org/10.1021/acs.jcim.3c01832>.
- ¹⁰¹ Vaissier, V., Sharma, S. C., Schaettle, K., Zhang, T. & Head-Gordon, T. Computational optimization of electric fields for improving catalysis of a designed kemp eliminase. *ACS Catalysis* **8**, 219–227 (2017). URL <http://dx.doi.org/10.1021/acscatal.7b03151>.
- ¹⁰² Liang, J., Feng, X., Hait, D. & Head-Gordon, M. Revisiting the performance of time-dependent density functional theory for electronic excitations: Assessment of 43 popular and recently developed functionals from rungs one to four. *J. Chem. Theory Comput.* **18**, 3460–3473 (2022).
- ¹⁰³ Spotte-Smith, E. W. C. *et al.* A database of molecular properties integrated in the materials project. *Digital Discovery* **2**, 1862–1882 (2023). URL <http://dx.doi.org/10.1039/D3DD000153A>.
- ¹⁰⁴ Seguin, T. J., Hahn, N. T., Zavadil, K. R. & Persson, K. A. Elucidating non-aqueous solvent stability and associated decomposition mechanisms for mg energy storage applications from first-principles. *Frontiers in Chemistry* **7** (2019). URL <http://dx.doi.org/10.3389/fchem.2019.00175>.
- ¹⁰⁵ Kim, D., King, D. S., Zhong, P. & Cheng, B. Learning charges and long-range interactions from energies and forces (2024). URL <https://arxiv.org/abs/2412.15455>. 2412.15455.
- ¹⁰⁶ Xiao, Y. *et al.* Understanding interface stability in solid-state batteries. *Nature Reviews Materials* **5**, 105–126 (2019). URL <http://dx.doi.org/10.1038/s41578-019-0157-5>.

- ¹⁰⁷ Rao, R. R. *et al.* Operando identification of site-dependent water oxidation activity on ruthenium dioxide single-crystal surfaces. *Nature Catalysis* **3**, 516–525 (2020). URL <http://dx.doi.org/10.1038/s41929-020-0457-6>.
- ¹⁰⁸ DaniFeri. Point defects in crystal structures: Interstitial and substitutional atom, vacancy, frenkel defect (2011).
- ¹⁰⁹ Batatia, I. *et al.* A foundation model for atomistic materials chemistry (2023). 2401.00096.
- ¹¹⁰ Vargas, S., Gee, W. & Alexandrova, A. High-throughput quantum theory of atoms in molecules (qtaim) for geometric deep learning of molecular and reaction properties. *Digit. Discov.* **3**, 987–998 (2024). URL <https://pubs.rsc.org/en/content/articlehtml/2024/dd/d4dd00057a>
<https://pubs.rsc.org/en/content/articlelanding/2024/dd/d4dd00057a>.
- ¹¹¹ Boiko, D. A., Reschützger, T., Sanchez-Lengeling, B., Blau, S. M. & Gomes, G. Advancing molecular machine (learned) representations with stereoelectronics-infused molecular graphs. *arXiv preprint arXiv:2408.04520* (2024). URL <https://arxiv.org/abs/2408.04520v1>.
- ¹¹² Dreuw, A. & Head-Gordon, M. Single-reference ab initio methods for the calculation of excited states of large molecules. *Chem. Rev.* **105**, 4009–4037 (2005).
- ¹¹³ Loos, P.-F., Scemama, A. & Jacquemin, D. The quest for highly accurate excitation energies: A computational perspective. *J. Phys. Chem. Lett.* **11**, 2374–2383 (2020).
- ¹¹⁴ Vargas, S., Hennefarth, M. R., Liu, Z. & Alexandrova, A. N. Machine learning to predict diels-alder reaction barriers from the reactant state electron density. *J. Chem. Theo. Comput.* **17**, 6203–6213 (2021). URL <https://pubs.acs.org/doi/full/10.1021/acs.jctc.1c00623>.
- ¹¹⁵ Li, S. C. *et al.* When do quantum mechanical descriptors help graph neural networks to predict chemical properties? *J. Am. Chem. Soc.* **146**, 23103–23120 (2024). URL <https://pubs.acs.org/doi/full/10.1021/jacs.4c04670>.
- ¹¹⁶ Najibi, A. & Goerigk, L. The nonlocal kernel in van der Waals density functionals as an additive correction: An extensive analysis with special emphasis on the B97M-V and ω B97M-V approaches. *J. Chem. Theory Comput.* **14**, 5725–5738 (2018).
- ¹¹⁷ Helgaker, T., Klopper, W. & Tew, D. P. Quantitative quantum chemistry. *Mol. Phys.* **106**, 2107–2143 (2008).
- ¹¹⁸ Raghavachari, K., Trucks, G. W., Pople, J. A. & Head-Gordon, M. A fifth-order perturbation comparison of electron correlation theories. *Chem. Phys. Lett.* **157**, 479 (1989).
- ¹¹⁹ Karton, A. Quantum mechanical thermochemical predictions 100 years after the schrödinger equation. *Ann. Rep. Comp. Chem.* **18**, 123–166 (2022).

- ¹²⁰ Smith, J. S. *et al.* Approaching coupled cluster accuracy with a general-purpose neural network potential through transfer learning. *Nat. Comm.* **2019 10:1** **10**, 1–8 (2019). URL <https://www.nature.com/articles/s41467-019-10827-4>.
- ¹²¹ Käser, S., Koner, D., Christensen, A. S., Lilienfeld, O. A. V. & Meuwly, M. Machine learning models of vibrating h2co: Comparing reproducing kernels, fchl, and physnet. *J. Phys. Chem. A* **124**, 8853–8865 (2020). URL <https://pubs.acs.org/doi/full/10.1021/acs.jpca.0c05979>.
- ¹²² Chen, M. S. *et al.* Data-efficient machine learning potentials from transfer learning of periodic correlated electronic structure methods: Liquid water at afqmc, ccSD, and ccSD(t) accuracy. *J. Chem. Theo. Comput.* **19**, 4510–4519 (2023). URL <https://pubs.acs.org/doi/full/10.1021/acs.jctc.2c01203>.
- ¹²³ Khazieva, E. O., Chtchelkatchev, N. M. & Ryltsev, R. E. Transfer learning for accurate description of atomic transport in al-cu melts. *J. Chem. Phys.* **161**, 174101 (2024). URL [/aip/jcp/article/161/17/174101/3318454/Transfer-learning-for-accurate-description-of](https://aip/jcp/article/161/17/174101/3318454/Transfer-learning-for-accurate-description-of).
- ¹²⁴ Witte, J., Neaton, J. B. & Head-Gordon, M. Push it to the limit: comparing periodic and local approaches to density functional theory for intermolecular interactions. *Mol. Phys.* **117**, 1298–1305 (2019).
- ¹²⁵ Bosoni, E. *et al.* How to verify the precision of density-functional-theory implementations via reproducible and universal workflows. *Nat. Rev. Phys.* **6**, 45–58 (2024).
- ¹²⁶ Karton, A. & De Oliveira, M. T. Good practices in database generation for benchmarking density functional theory. *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **15**, e1737 (2025).
- ¹²⁷ Spiekermann, K. A., Pattanaik, L. & Green, W. H. High Accuracy Barrier Heights, Enthalpies, and Rate Coefficients for Chemical Reactions. *Sci. Data* **9**, 417 (2022).
- ¹²⁸ Liu, X., Spiekermann, K. A., Lagnajit, P., Green, W. & Head-Gordon, M. Revisiting a large and diverse data set for barrier heights and reaction energies: Best practices in density functional theory calculations for chemical kinetics. (*in preparation*) (2025).
- ¹²⁹ Shee, J., Loipersberger, M., Hait, D., Lee, J. & Head-Gordon, M. Revealing the nature of electron correlation in transition metal complexes with symmetry breaking and chemical intuition. *J. Chem. Phys.* **154**, 194109 (2021).
- ¹³⁰ Zhang, I. Y. & Grüneis, A. Coupled cluster theory in materials science. *Frontiers in Materials* **6**, 123 (2019).
- ¹³¹ Lejaeghere, K. *et al.* Reproducibility in density functional theory calculations of solids. *Science* **351**, aad3000 (2016).
- ¹³² Abdelmaqsoud, K., Shuaibi, M., Kolluru, A., Cheula, R. & Kitchin, J. R. Investigating the error imbalance of large-scale machine learning potentials in catalysis. *Catalysis Science & Technology* **14**, 5899–5908 (2024). URL <http://dx.doi.org/10.1039/D4CY00615A>.

- ¹³³ Quiton, S. J., Wu, H., Xing, X., Lin, L. & Head-Gordon, M. The staggered mesh method: Accurate exact exchange toward the thermodynamic limit for solids. *J. Chem. Theory Comput.* **20**, 7958–7968 (2024).
- ¹³⁴ Borlido, P., Doumont, J., Tran, F., Marques, M. A. & Botti, S. Validation of pseudopotential calculations for the electronic band gap of solids. *J. Chem. Theory Comput.* **16**, 3620–3627 (2020).
- ¹³⁵ Rossomme, E. *et al.* The good, the bad, and the ugly: Pseudopotential inconsistency errors in molecular applications of density functional theory. *J. Chem. Theory Comput.* **19**, 2827–2841 (2023).
- ¹³⁶ Li, W.-L., Chen, K., Rossomme, E., Head-Gordon, M. & Head-Gordon, T. Greater transferability and accuracy of norm-conserving pseudopotentials using nonlinear core corrections. *Chem. Sci.* **14**, 10934–10943 (2023). URL <http://dx.doi.org/10.1039/D3SC03709F>.
- ¹³⁷ Van Voorhis, T. & Head-Gordon, M. A geometric approach to direct minimization. *Mol. Phys.* **100**, 1713 (2002).
- ¹³⁸ Liakos, D. G. & Neese, F. Is it possible to obtain coupled cluster quality energies at near density functional theory cost? domain-based local pair natural orbital coupled cluster vs modern density functional theory. *J. Chem. Theory Comput.* **11**, 4054–4063 (2015).
- ¹³⁹ Liakos, D. G., Sparta, M., Kesharwani, M. K., Martin, J. M. & Neese, F. Exploring the accuracy limits of local pair natural orbital coupled-cluster theory. *J. Chem. Theory Comput.* **11**, 1525–1539 (2015).
- ¹⁴⁰ Santra, G. & Martin, J. M. Performance of localized-orbital coupled-cluster approaches for the conformational energies of longer n-alkane chains. *J. Phys. Chem. A* **126**, 9375–9391 (2022).
- ¹⁴¹ Gray, M. & Herbert, J. M. Assessing the domain-based local pair natural orbital (dlpno) approximation for non-covalent interactions in sizable supramolecular complexes. *J. Chem. Phys.* **161** (2024).
- ¹⁴² Wang, Z. *et al.* Local second-order moller–plesset theory with a single threshold using orthogonal virtual orbitals: Theory, implementation, and assessment. *J. Chem. Theory Comput.* **19**, 7577–7591 (2023).
- ¹⁴³ Shi, T. *et al.* Local second order moller–plesset theory with a single threshold using orthogonal virtual orbitals: A distributed memory implementation. *J. Chem. Theory Comput.* **20**, 8010–8023 (2024).
- ¹⁴⁴ Garrison, A. G. *et al.* Applying large graph neural networks to predict transition metal complex energies using the tmqm_wb97mv data set. *Journal of Chemical Information and Modeling* **63**, 7642–7654 (2023). URL <https://doi.org/10.1021/acs.jcim.3c01226>. PMID: 38049389, <https://doi.org/10.1021/acs.jcim.3c01226>.

- ¹⁴⁵ Balcells, D. & Skjelstad, B. B. tmqm dataset—quantum geometries and properties of 86k transition metal complexes. *Journal of Chemical Information and Modeling* **60**, 6135–6146 (2020). URL <https://doi.org/10.1021/acs.jcim.0c01041>. PMID: 33166143, <https://doi.org/10.1021/acs.jcim.0c01041>.
- ¹⁴⁶ Gong, S. *et al.* Bamboo: a predictive and transferable machine learning force field framework for liquid electrolyte development. *arXiv preprint arXiv:2404.07181* (2024). URL <https://arxiv.org/abs/2404.07181v4>.
- ¹⁴⁷ Hu, W. *et al.* Ogb-lsc: A large-scale challenge for machine learning on graphs. *Advances in Neural Information Processing Systems (NeurIPs)* (2021). URL <https://arxiv.org/abs/2103.09430v3>.
- ¹⁴⁸ Nakata, M., Shimazaki, T., Hashimoto, M. & Maeda, T. Pubchemqc pm6: Data sets of 221 million molecules with optimized molecular geometries and electronic properties. *Journal of Chemical Information and Modeling* **60**, 5891–5899 (2020). URL <https://doi.org/10.1021/acs.jcim.0c00740>.
- ¹⁴⁹ Nakata, M. & Shimazaki, T. Pubchemqc project: A large-scale first-principles electronic structure database for data-driven chemistry. *J. Chem. Info. Model.* **57**, 1300–1308 (2017). URL <https://pubs.acs.org/doi/full/10.1021/acs.jcim.7b00083>.
- ¹⁵⁰ Zaidi, S. *et al.* Pre-training via denoising for molecular property prediction. *11th International Conference on Learning Representations, ICLR 2023* (2022). URL <https://arxiv.org/abs/2206.00133v2>.
- ¹⁵¹ Liao, Y.-L., Smidt, T., Shuaibi, M. & Das, A. Generalizing denoising to non-equilibrium structures improves equivariant force fields (2024). URL <https://arxiv.org/abs/2403.09549>. 2403.09549.
- ¹⁵² Gardner, J. L. A., Baker, K. T. & Deringer, V. L. Synthetic pre-training for neural-network interatomic potentials. *Mach. Learn.: Sci. Technol.* **5**, 015003 (2024). URL <https://iopscience.iop.org/article/10.1088/2632-2153/ad1626><https://iopscience.iop.org/article/10.1088/2632-2153/ad1626/meta>.
- ¹⁵³ Magar, R., Wang, Y. & Farimani, A. B. Crystal twins: self-supervised learning for crystalline material property prediction. *npj Comput. Mater.* **2022 8:1 8**, 1–8 (2022). URL <https://www.nature.com/articles/s41524-022-00921-5>.
- ¹⁵⁴ Sakai, Y. *et al.* Self-supervised learning with atom replacement for catalyst energy prediction by graph neural networks. *Proc. Comp. Sci.* **222**, 458–467 (2023).
- ¹⁵⁵ Kovács, D. P. *et al.* Mace-off: Transferable short range machine learning force fields for organic molecules. *arXiv preprint arXiv:2312.15211* (2023). URL <https://arxiv.org/abs/2312.15211v3>.
- ¹⁵⁶ Shoghi, N. *et al.* From molecules to materials: Pre-training large generalizable models for atomic property prediction. *12th International Conference on Learning Representations, ICLR 2024* (2023). URL <https://arxiv.org/abs/2310.16802v2>.

- ¹⁵⁷ Pasini, M. L. *et al.* Scalable training of trustworthy and energy-efficient predictive graph foundation models for atomistic materials modeling: A case study with hydragnn. *arXiv preprint arXiv:2406.12909* (2024). URL <https://arxiv.org/abs/2406.12909v4>.
- ¹⁵⁸ Chmiela, S. *et al.* Machine learning of accurate energy-conserving molecular force fields. *Sci. Adv.* **3** (2017). URL <https://www.science.org/doi/10.1126/sciadv.1603015>.
- ¹⁵⁹ Pan, J. Can machines learn with hard constraints? *Nat. Comp. Sci. 2021 1:4* **1**, 244–244 (2021). URL <https://www.nature.com/articles/s43588-021-00067-2>.
- ¹⁶⁰ Amin, I., Raja, S. & Krishnapriyan, A. S. Towards fast, specialized machine learning force fields: Distilling foundation models via energy Hessians. In *The Thirteenth International Conference on Learning Representations* (2025). URL <https://openreview.net/forum?id=1durmugh3I>.
- ¹⁶¹ Kelvinius, F. E., Georgiev, D., Toshev, A. P. & Gasteiger, J. Accelerating molecular graph neural networks via knowledge distillation. *Advances in Neural Information Processing Systems* **36** (2023). URL <https://arxiv.org/abs/2306.14818v2>.
- ¹⁶² Wang, W., Axelrod, S. & Gómez-Bombarelli, R. Differentiable molecular simulations for control and learning. In *ICLR 2020 Workshop on Integration of Deep Neural Models and Differential Equations* (2019).
- ¹⁶³ Thaler, S. & Zavadlav, J. Learning neural network potentials from experimental data via differentiable trajectory reweighting. *Nat. Comm. 2021 12:1* **12**, 1–10 (2021). URL <https://www.nature.com/articles/s41467-021-27241-4>.
- ¹⁶⁴ Šípka, M., Dietschreit, J. C., Grajciar, L. & Gómez-Bombarelli, R. Differentiable simulations for enhanced sampling of rare events. *Proc. Mach. Learn. Res.* **202**, 31990–32007 (2023). URL <https://arxiv.org/abs/2301.03480v2>.
- ¹⁶⁵ Navarro, C., Majewski, M. & Fabritiis, G. D. Top-down machine learning of coarse-grained protein force fields. *J. Chem. Theo. Comput.* **19**, 7518–7526 (2023). URL <https://pubs.acs.org/doi/full/10.1021/acs.jctc.3c00638>.
- ¹⁶⁶ Gangan, A. S. *et al.* Force field optimization by end-to-end differentiable atomistic simulation. *arXiv preprint arXiv:2409.13844* (2024). URL <https://arxiv.org/abs/2409.13844v1>.
- ¹⁶⁷ Blondel, M. *et al.* Efficient and modular implicit differentiation. In *Advances in Neural Information Processing Systems (NeurIPS)* (Google Research, 2022).
- ¹⁶⁸ Chen, R. T. Q., Rubanova, Y., Bettencourt, J. & Duvenaud, D. K. Neural ordinary differential equations. In Bengio, S. *et al.* (eds.) *Advances in Neural Information Processing Systems*, vol. 31 (Curran Associates, Inc., 2018). URL https://proceedings.neurips.cc/paper_files/paper/2018/file/69386f6bb1dfed68692a24c8686939b9-Paper.pdf.

- ¹⁶⁹ Bolhuis, P. G., Brotzakis, Z. F. & Keller, B. G. Optimizing molecular potential models by imposing kinetic constraints with path reweighting. *J. Chem. Phys.* **159**, 074102 (2023). URL <https://doi.org/10.1063/5.0151166>.
- ¹⁷⁰ Wang, X., Zhao, H., Tu, W. & Yao, Q. Automated 3d pre-training for molecular property prediction. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* **23**, 2419–2430 (2023). URL <http://arxiv.org/abs/2306.07812><http://dx.doi.org/10.1145/3580305.3599252>.
- ¹⁷¹ Zuo, Y. *et al.* Performance and cost assessment of machine learning interatomic potentials. *J. Phys. Chem. A* **124**, 731–745 (2020). URL <https://pubs.acs.org/doi/10.1021/acs.jpca.9b08723>.
- ¹⁷² Liu, Y., He, X. & Mo, Y. Discrepancies and error evaluation metrics for machine learning interatomic potentials. *npj Comput. Mater.* **9**, 174 (2023). URL <http://arxiv.org/abs/2306.11639><https://www.nature.com/articles/s41524-023-01123-3>.
- ¹⁷³ Fu, X. *et al.* Forces are not enough: Benchmark and critical evaluation for machine learning force fields with molecular simulations. *Transactions on Machine Learning Research* (2023). URL <https://openreview.net/forum?id=A8pqQipwkt>. Survey Certification.
- ¹⁷⁴ Kapil, V. *et al.* The first-principles phase diagram of monolayer nanoconfined water. *Nature* **609**, 512–516 (2022).
- ¹⁷⁵ Wang, H. *et al.* Evaluating self-supervised learning for molecular graph embeddings. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track* (2022). URL <http://arxiv.org/abs/2206.08005>.
- ¹⁷⁶ Rohskopf, A. *et al.* Exploring model complexity in machine learned potentials for simulated properties. *J. Mater. Res.* **38**, 5136–5150 (2023). URL <https://link.springer.com/10.1557/s43578-023-01152-0>.
- ¹⁷⁷ Liu, Y. & Mo, Y. Learning from models: high-dimensional analyses on the performance of machine learning interatomic potentials. *npj Comput. Mater.* **10** (2024).
- ¹⁷⁸ Lakshminarayanan, B., Pritzel, A. & Blundell, C. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in Neural Information Processing Systems* **2017-December**, 6403–6414 (2016). URL <https://arxiv.org/abs/1612.01474v3>.
- ¹⁷⁹ Peterson, A. A., Christensen, R. & Khorshidi, A. Addressing uncertainty in atomistic machine learning. *Phys. Chem. Chem. Phys.* **19**, 10978–10985 (2017). URL <https://pubs.rsc.org/en/content/articlehtml/2017/cp/c7cp00375g><https://pubs.rsc.org/en/content/articlelanding/2017/cp/c7cp00375g>.
- ¹⁸⁰ Imbalzano, G. *et al.* Uncertainty estimation for molecular dynamics and sampling. *J. Chem. Phys.* **154**, 74102 (2021). URL [/aip/jcp/article/154/7/074102/200970/Uncertainty-estimation-for-molecular-dynamics-and](http://aip/jcp/article/154/7/074102/200970/Uncertainty-estimation-for-molecular-dynamics-and).

- ¹⁸¹ Morrow, J. D., Gardner, J. L. & Deringer, V. L. How to validate machine-learned interatomic potentials. *J. Chem. Phys.* **158**, 121501 (2023). URL [/aip/jcp/article/158/12/121501/2881528/How-to-validate-machine-learned-interatomic](#).
- ¹⁸² Bihani, V. *et al.* Egraffbench: evaluation of equivariant graph neural network force fields for atomistic simulations. *Digit. Discov.* **3**, 759–768 (2024). URL <https://pubs.rsc.org/en/content/articlehtml/2024/dd/d4dd00027g><https://pubs.rsc.org/en/content/articlelanding/2024/dd/d4dd00027g>.
- ¹⁸³ Rowan benchmarks. URL <https://benchmarks.rowansci.com/>.
- ¹⁸⁴ Chiang, Y. *et al.* MLIP arena: Advancing fairness and transparency in machine learning interatomic potentials through an open and accessible benchmark platform. In *AI for Accelerated Materials Design - ICLR 2025* (2025). URL <https://openreview.net/forum?id=ysKfIavYQE>.
- ¹⁸⁵ Wood, B. M. *et al.* Uma: A family of universal models for atoms (2025). URL <https://ai.meta.com/research/publications/uma-a-family-of-universal-models-for-atoms/>.
- ¹⁸⁶ Fu, X. *et al.* Learning smooth and expressive interatomic potentials for physical property prediction (2025). URL <https://arxiv.org/abs/2502.12147>. 2502.12147.
- ¹⁸⁷ Goerigk, L. *et al.* A look at the density functional theory zoo with the advanced gmtkn55 database for general main group thermochemistry, kinetics and noncovalent interactions. *Physical Chemistry Chemical Physics* **19**, 32184–32215 (2017). URL <http://dx.doi.org/10.1039/C7CP04913G>.
- ¹⁸⁸ Abadi, M. *et al.* TensorFlow: Large-scale machine learning on heterogeneous systems (2015). URL <http://tensorflow.org/>. Software available from tensorflow.org.
- ¹⁸⁹ Chollet, F. *et al.* Keras (2015). URL <https://github.com/fchollet/keras>.
- ¹⁹⁰ Paszke, A. *et al.* Automatic differentiation in pytorch. In *NIPS-W* (2017).
- ¹⁹¹ OpenAI. Chatgpt (2023). URL <https://chat.openai.com/chat>.
- ¹⁹² Muller, J. & Lewis, G. The ethics of emerging technology: The era of artificial intelligence—dr. teresa head-gordon (2024). URL <https://www.buzzsprout.com/1204337/episodes/15285470-the-ethics-of-emerging-technology-the-era-of-artificial-intelligence-dr-teresa-head-gordon>.
- ¹⁹³ Cook-Deegan, R. M. *The Gene Wars: Science, Politics, and the Human Genome* (W.W. Norton & Company, 1994). URL <https://books.google.com/books?id=2shp9MyPaPUC>.

5 ACKNOWLEDGMENTS

The authors thank Eric Qu, Yingze Wang, Joseph Cavanagh, Kunyang O. Sun, Kareem Hegazy, Joseph P. Heindel, Peichen Zhong, Anup Kumar, Senwei Liang, Selim Sami, R. Allen LaCour, Sruthy K. Chandy, Ruoqi Zhao, and Dorian Bagni for their contributions in discussions on August 6, 2024 at UC Berkeley. T.H.-G. thanks the CPIMS program, Office of Science, Office of Basic Energy Sciences, Chemical Sciences Division of the U.S. Department of Energy under Contract DE-AC02-05CH11231 for support of the machine learning research. S.M.B. acknowledges support by the Center for High Precision Patterning Science (CHiPPS), an Energy Frontier Research Center funded by the U.S. Department of Energy, Office of Science, Basic Energy Sciences at Lawrence Berkeley National Laboratory under contract DE-AC02-05CH11231. S.V. was supported by the DOE’s National Nuclear Security Administration’s Office of Defense Nuclear Nonproliferation Research and Development (NA-22) as part of the NextGen Nonproliferation Leadership Development Program. This work was also supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing, and Office of Basic Energy Sciences, via the Scientific Discovery through Advanced Computing (SciDAC) program for the mathematical and computing foundations of machine learning models. This work was also supported by the U.S. Department of Energy, Office of Science, Energy Earthshot initiatives as part of the Center for Ionomer-based Water Electrolysis at Lawrence Berkeley National Laboratory under Award Number DE-AC02-05CH11231, as well as Toyota Research Institute as part of the Synthesis Advanced Research Challenge. This work used computational resources provided by the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility operated under Contract DE-AC02-05CH11231.


6 AUTHOR CONTRIBUTIONS STATEMENT

S.M.B., B.C., A.K., and T.H.-G. defined the topics and scope of the perspective. E.C.-Y.Y., Y.L., J.C., S.R., T.K., S.V., W.X. made the figures. E.C.-Y.Y., Y.L., J.C., S.R., T.K., S.V., P.Z., W.X., M.H.-G., S.M.B., B.C., A.K., and T.H.-G. wrote the paper. All authors discussed the results and made comments and edits to the manuscript.

7 COMPETING INTERESTS STATEMENT

The authors declare no competing interests.

8 Graphical abstract



9 Short Summary

We examine the historical development and underlying principles of foundation models realized in language and vision, and propose how physics-infused machine learning interaction potentials could dramatically transform at scale to create transformative foundation models for chemistry and materials science.