Accelerating Machine Learning Force Field Inference with Atom Type Specific Scheduling

Yuko Kinoshita[1], Meguru Yamazaki[2], Yuta Yoshimoto[2], Shintaro Izumi[1], Atsuki Inoue[1], Hiroshi Kawaguchi[1], Yasufumi Sakai[2]

[1] Graduate School of Science, Technology and Innovation, Kobe University, Japan [2]Fujitsu Limited, Japan

Corresponding author: kinoshita.yuko@cs28.cs.kobe-u.ac.jp

Abstract

Molecular dynamics (MD) simulations using machine learning force fields (MLFFs) have enabled high-accuracy modeling of complex materials systems. However, the significant computational cost of MLFF-based MD remains a challenge, especially for large-scale simulations required in materials discovery. We propose an efficient MD simulation method that adaptively reduces the MLFF inference frequency for atom types with smaller displacements, thereby accelerating simulations without compromising accuracy. We implement the proposed approach in DeePMD and evaluate it on crystalline TiO_2 anatase with 6,144 atoms. Our experiments demonstrate that the proposed method achieves approximately 1.16× speedup compared to conventional DeePMD, while preserving the accuracy of key physical properties such as the radial distribution function, temperature, and density. This atom-type specific inference scheduling provides a practical pathway to scalable, resource-efficient MD simulations for materials design. The proposed method is also expected to be effective for future applications in large-scale, non-periodic systems such as amorphous membranes.

Introduction

Molecular dynamics (MD) simulations are essential in materials science, but large-scale and long- timescale simulations remain challenging due to the limitations of classical force fields and the computational expense of quantum calculations (Behler 2015). Recent progress in machine learning force fields (MLFFs), such as DeePMD (Zhang et al. 2018a) and graph neural network approaches, has enabled near-first-principles accuracy for a wide range of systems (Zhang et al. 2018b; Schütt et al. 2017). These advances include general-purpose potential energy surfaces (Bartók et al. 2010), neural message passing frameworks for quantum chemistry (Gilmer et al. 2017), improved accuracy for molecular dynamics simulations with machine-learned force fields (Chmiela et al. 2018), automatic selection of atomic fingerprints and reference configurations for ML potentials (Imbalzano et al. 2018), and scalable parallel GNN algorithms for interatomic potentials (Park et al. 2024). However, their high inference cost, especially for large systems, is a significant bottleneck (Zeng et al. 2023). Although parallelization and multiple-time-step methods have improved efficiency (Fu et al. 2023; Ferrarotti et al. 2015), existing approaches treat all atoms equivalently, regardless of their mobility or physical relevance.

Here, we propose an efficient strategy that adaptively reduces inference frequency for atom types with smaller displacements while maintaining accuracy for key physical observables. We implement and validate this approach in DeePMD for TiO_2 systems (Calegari Andrade and Selloni 2020), chosen for their distinct Ti and O displacement characteristics and importance in solid-state and catalytic applications requiring large- scale, long-timescale simulations, demonstrating its utility for scalable, resource-efficient MD in material discovery.

Methods

Atom-Type Specific Inference Scheduling

As illustrated in Figure 1, we propose a scheduling scheme to reduce the frequency of MLFF inferences for atom types with small displacements. Specifically, at every N-th time step, the forces for all atoms are calculated using MLFF inference. For the other steps, only the forces of the mobile atoms (e.g., O in TiO_2, which exhibit larger displacements) are updated via MLFF inference, while immobile atoms retain their last force corrections computed at the most recent full inference step to maintain accuracy. Let A_{mobile} be the set of atom types for which high-frequency inference is performed. Let N be the scheduling period, and let ΔF_i denote the last stored force correction for atom i (computed during the final full inference step). Subsequently, the force on atom i at time step t is calculated as follows:

$$\boldsymbol{F}_i^t = \begin{cases} \boldsymbol{F}_i \ (i \in A_{all}) & if\ t\ mod\ N = 0 \\ \boldsymbol{F}_i' \ (i \in A_{mobile}) + \Delta \boldsymbol{F}_i & if\ t\ mod\ N \neq 0 \end{cases} \quad (1)$$

$$\Delta \boldsymbol{F}_i = \boldsymbol{F}_i \ (i \in A_{all}) - \boldsymbol{F}_i' \ (i \in A_{mobile}) \quad if\ t\ mod\ N = 0 \quad (2)$$

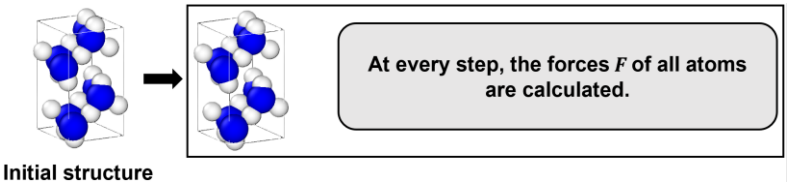

(a) traditional method

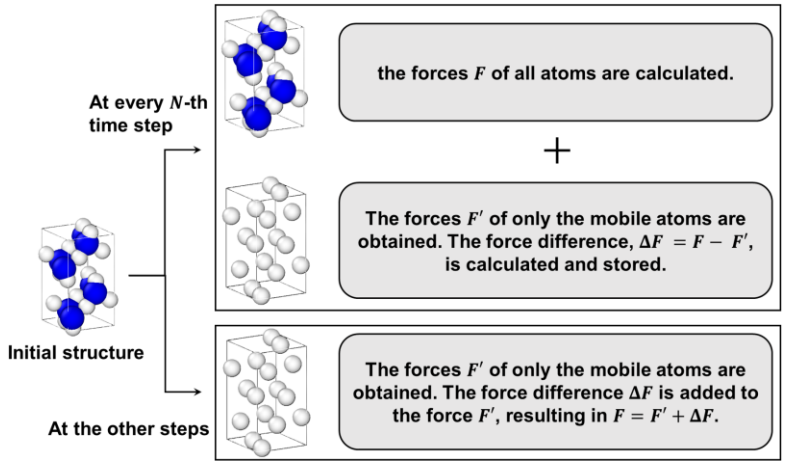

(b) proposed method

Figure 1: Workflow of atom-type specific inference scheduling. At every *N*-th time step, full model inference is performed for all atoms. At the other steps, inference is performed only for the mobile atoms, with stored force corrections added for the immobile atoms.

Suppose that the fraction of timesteps with full inference is $f_{all} = 1/N$ and the fraction of atoms requiring high-frequency inference is p_{mobile}. Thus, the total computational cost over N steps is:

$$C = f_{all} \times 1 + (1 - f_{all}) \times p_{mobile} \quad (3)$$

The theoretical speedup is denoted as $1/C$. For example, for TiO_2 ($p_{highfreq} = 2/3$ for oxygen atoms and $N = 10$), the expected speedup is approximately 1.30×.

Experimental Setup

We evaluated the proposed atom-type specific inference scheduling using molecular-dynamics (MD) simulations of anatase TiO_2. DeePMD models were trained using a public dataset containing 12 TiO_2 crystal structures and 46,692 configurations (Zeng et al. 2023). For the MD simulations, anatase TiO_2 structures (6,144 atoms) were used under the NPT ensemble at 300 K with a time step of 0.1 fs for 10,000 steps; furthermore, the simulations were performed using the LAMMPS package (Thompson et al. 2022) in conjunction with the DeePMD-kit plugin. To assess the effect of atom type specific scheduling, we performed simulations with N = 1, 2, 4, 8, and 10, where N is the update interval for non-target atom types (e.g., Ti atoms in TiO_2).

For the baseline, MLFF inference was performed for all atoms at every step. In the proposed method, only oxygen atoms were inferred at every step, whereas titanium atoms were updated at every N steps using stored force corrections. All simulations used the Nose-Hoover thermostat for temperature control. We compared the accuracy (radial distribution function and atomic forces) and computational efficiency (timesteps per second) of the two methods.

The proposed method was integrated into the DeePMD-kit MD engine using the TensorFlow framework (Pang et al. 2020). At each MD integration step, the scheduling mechanism determines whether to perform full or partial inference while maintaining a registry of mobile atoms (A_{mobile}) and storing force corrections ($\Delta \boldsymbol{F_i}$) for immobile atoms. This approach is compatible with GPU-accelerated calculations and can be easily incorporated into standard simulation workflows.

Results and Discussion

Accuracy of Physical Properties

Figure 2 shows a comparison of the radial distribution functions (RDFs) for Ti-Ti, Ti-O, and O-O pairs obtained from the baseline DeePMD and the proposed method for anatase TiO_2 at 300 K. We compared the RDF profiles obtained for N = 1(baseline), 2, 4, 8, 10 in our scheduling scheme, as well as with reference structural data (Calegari Andrade and Selloni 2020). The RDF profiles for the proposed method at each N closely match those of the baseline, with peak positions and heights preserved, indicating that the crystal structure is maintained. Temperature and density remained stable after the first 2,000 steps, fluctuating around 300 K and 3.82–3.84 g/cm^3, confirming thermodynamic stability.

Table 1 further quantifies this agreement by reporting the Pearson correlation coefficients between the RDF profiles produced at various inference intervals and the baseline. Even when the inference interval for all atoms is extended up to 10 steps, the correlation remains exceptionally high (all > 0.99), demonstrating that the proposed scheduling strategy maintains structural accuracy across a range of step intervals.

Table 1: Pearson Correlation with Baseline RDF Across Inference Steps

Inference interval for all atoms	Pearson correlation coefficient
2 steps	0.9968
4 steps	0.9951
8 steps	0.9924
10 steps	0.9923

Force Accuracy

Figure 3 shows the time evolution of force components (F_x, F_y, F_z) for representative Ti and O atoms, respectively, as obtained from both the baseline and the proposed method under intervals N = 10. For all force components, the results from the proposed method closely matched those from the baseline, especially during the initial 2,000 timesteps. At later timesteps, small deviations appear for components with large force fluctuations (e.g., F_y for O atom), primarily due to error accumulation from stored force corrections in non-target atom types. Overall, the force accuracy for Ti and O atoms remained high for N=10, and these deviations were not found to significantly affect the crystal structure or thermodynamic stability.

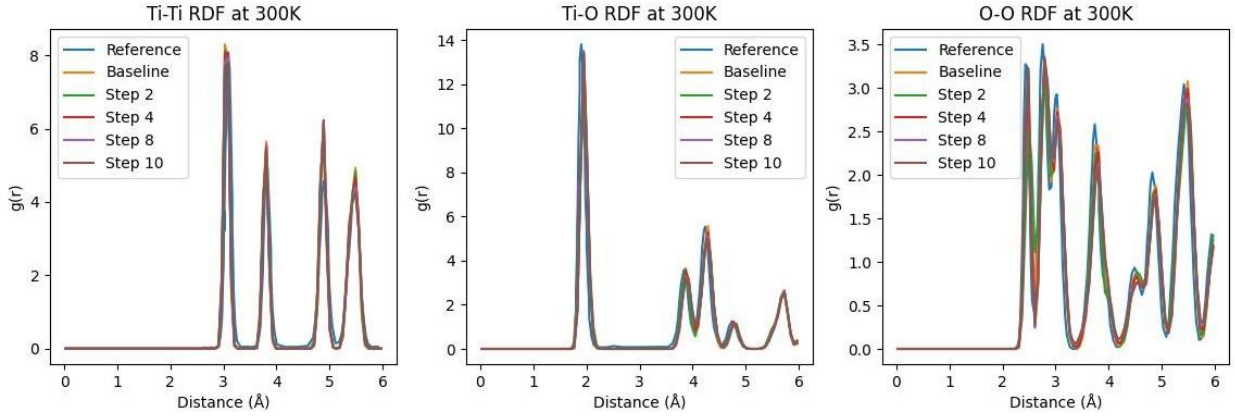

Figure 2: Comparison of RDFs for Ti-Ti, Ti-O, and O-O pairs from baseline and proposed method for anatase TiO_2 at 300 K.

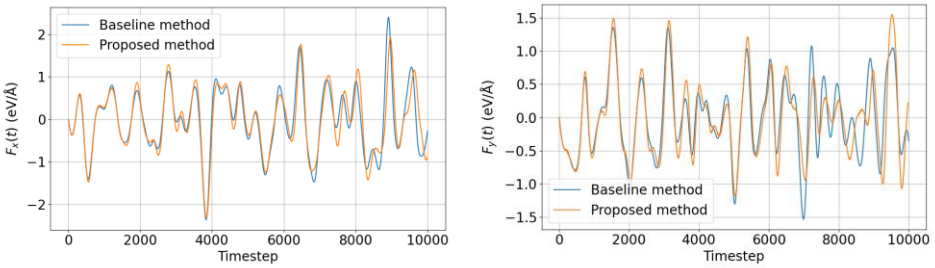

(a) F_x for Ti atom (b) F_y for Ti atom

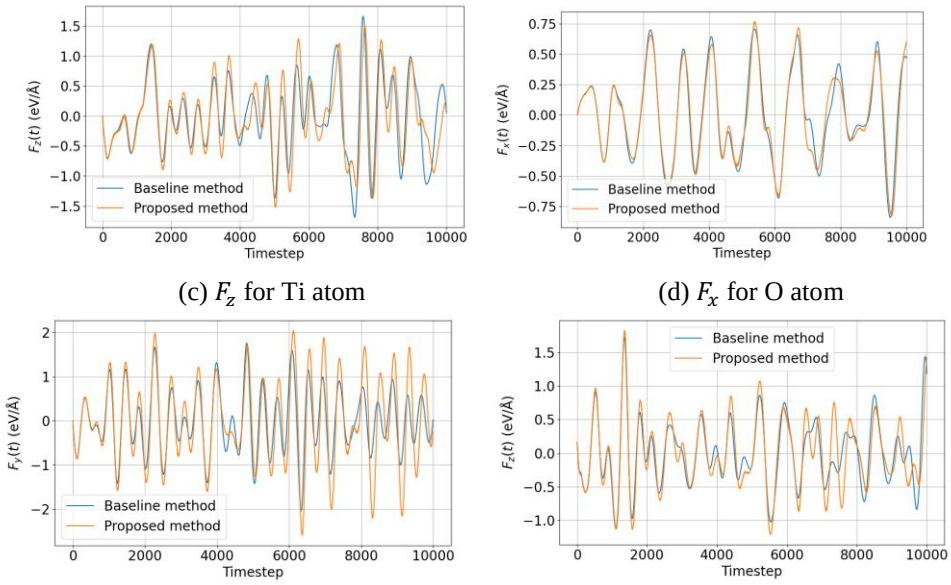

(c) F_z for Ti atom (d) F_x for O atom

(e) F_y for O atom (f) F_z for O atom

Figure 3: Time evolution of force components for representative Ti and O atoms. The proposed method (which incorporates the stored force corrections ΔF) is compared with the baseline method for each component.

Computational Efficiency

Computational efficiency was assessed over three independent trials (Table 2). The baseline method achieved 31.3 timesteps per second, whereas the proposed method reached 36.2 timesteps per second, yielding a speedup of 1.16×. This performance agrees well with the theoretical estimate of 1.30×, which was derived from the reduction in inference calls; a slight performance loss was observed, which is attributed to the additional overhead required for force-correction computations. The reported computational efficiency reflects the case for N=10 in our scheme.

Table 2: Average MD performance over three trials.

Method	Timesteps/s	Speedup
Baseline (naive DeePMD)	31.3	1.00
Proposed	36.2	1.16
Theory(est.)	—	1.30

Discussion

In this study, we evaluated the effect of varying the MLFF inference frequency for all atoms in the TiO_2 anatase system using intervals of 1, 2, 4, 8, and 10 steps. For this system with small atomic displacements and limited mobility, performing inference every 10 steps maintained accuracy for key physical properties (structure, force, density, temperature) while improving computational efficiency by about 16% over the baseline model performing inference for all atoms at every step. The approach was validated only for TiO_2 systems with low atomic mobility. Its applicability to other systems—particularly those with higher mobility or longer simulation timescales—remains to be verified. The optimal inference interval likely depends on system and simulation conditions; thus, the favorable result at a 10-step interval should be regarded as case-specific.

Future studies should test this method across diverse materials and timescales, particularly in simulating relaxation processes, to examine error accumulation and optimize inference scheduling. At this stage, this approach can be viewed as an efficient but system-dependent method demonstrated to work effectively for solid-state systems.

Conclusion

In this study, we proposed an efficient MD simulation method using MLFFs by reducing the inference frequency for atom types with small displacements. The proposed method, implemented in DeePMD, maintained high accuracy for key physical properties (such as the radial distribution function, temperature, and density) while improving the computational efficiency by approximately 16% compared with the baseline model that performed MLFF inference for all atoms at every step. These results confirm that atom-type-specific inference scheduling can accelerate simulations without accuracy loss for systems with limited atomic mobility, such as crystalline TiO_2. Further validation on other materials and simulation conditions will be required to assess the general applicability and scalability of this approach.

References

Bartók, A. P., M. C. Payne, R. Kondor, and G. Csányi, "Machine Learning a General-Purpose Potential Energy Surface for High-Dimensional Materials Simulations," Physical Review Letters, Vol. 104, No. 13, p. 136403, Apr. 2010.

Behler, J., "First Principles Neural Network Potentials for Reactive Simulations of Large Molecular and Condensed Systems," International Journal of Quantum Chemistry, Vol. 115, No. 16, pp. 1032–1042, Aug. 2015.

Calegari Andrade, M. F., and A. Selloni, "Structure of disordered TiO2 phases from ab initio based deep neural network simulations," Physical Review Materials, Vol. 4, No. 11, 2020.

Chmiela, S., H. E. Sauceda, K. R. Müller, and A. Tkatchenko, "Towards exact molecular dynamics simulations with machine-learned force fields," Nature Communications, Vol. 9, No. 1, p. 3887, 2018.

Ferrarotti, M. J., S. Bottaro, A. Pérez-Villa, and G. Bussi, "Accurate Multiple Time Step in Biased Molecular Simulations," Journal of Chemical Theory and Computation, Vol. 11, No. 1, pp. 139-146, Jan. 2015.

Fu, X., A. Musaelian, A. Johansson, T. Jaakkola, and B. Kozinsky, "Learning Interatomic Potentials at Multiple Scales," arXiv:2310.13756, 2023.

Gilmer, J., S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl, "Neural Message Passing for Quantum Chemistry," Proceedings of the 34th International Conference on Machine Learning (ICML), Vol. 70, pp. 1263–1272, Aug. 2017.

Imbalzano, G., A. Anelli, D. Giofré, S. Klees, J. Behler, and M. Ceriotti, "Automatic selection of atomic fingerprints and reference configurations for machine-learning potentials," The Journal of Chemical Physics, Vol. 148, No. 24, p. 241730, 2018.

Pang, B., E. Nijkamp, and Y. N. Wu, "Deep Learning With TensorFlow: A Review," Journal of Educational and Behavioral Statistics, Vol. 45, No. 2, pp. 227–248, 2020.

Park, Y., J. Kim, S. Hwang, and S. Han, "Scalable Parallel Algorithm for Graph Neural Network Interatomic Potentials in Molecular Dynamics Simulations," Journal of Chemical Theory and Computation, 2024.

Schütt, K. T., F. Arbabzadah, S. Chmiela, K. R. Müller, and A. Tkatchenko, "SchNet: A Continuous-Filter Convolutional Neural Network for Modeling Quantum Interactions," Advances in Neural Information Processing Systems, Vol. 30, pp. 992–1002, Dec. 2017.

Thompson, A. P., H. M. Aktulga, R. Berger, D. S. Bolintineanu, W. M. Brown, P. S. Crozier, P. J. in 't Veld, A. Kohlmeyer, S. G. Moore, T. D. Nguyen, R. Shan, M. J. Stevens, J. Tranchida, C. Trott, and S. J. Plimpton, "LAMMPS - a flexible simulation tool for particle-based materials modeling at the atomic, meso, and continuum scales," Computer Physics Communications, Vol. 271, 2022.

Zeng, J., D. Zhang, D. Lu, P. Mo, Z. Li, Y. Chen, M. Rynik, L. Huang, Z. Li, S. Shi, Y. Wang, H. Ye, P. Tuo, J. Yang, Y. Ding, Y. Li, D. Tisi, Q. Zeng, H. Bao, and H. Wang, "DeePMD-kit v2: A software package for deep potential models," Journal of Chemical Physics, Vol. 159, No. 5, 2023.

Zhang, L., J. Han, H. Wang, R. Car, and W. E, "Deep Potential Molecular Dynamics: A Scalable Model with the Accuracy of Quantum Mechanics," Physical Review Letters, Vol. 120, No. 14, p. 143001, Apr. 2018.

Zhang, L., J. Han, H. Wang, W. A. Saidi, R. Car, and W. E., "End-to-end Symmetry Preserving Inter-atomic Potential Energy Model for Finite and Extended Systems," arXiv:1805.09003, 2018