

# Acceleration of First-Principles Atomistic Simulations with Bayesian Neural Networks

Tim Rensmeyer\*, Oliver Niggemann

*Professorship Computer Science in Mechanical Engineering  
Helmut Schmidt University/University of the Bundeswehr  
Hamburg, Germany  
\*rensmeyt@hsu-hh.de*

**Abstract** – Molecular dynamics simulations with first-principles methods, such as density functional theory, are a cornerstone in the development of new battery and fuel cell materials. However, due to their high computational demand, their application is mostly limited to small systems and short time horizons. AI-based methods are a promising approach for accelerating first-principles simulations while maintaining high simulation accuracy. A key challenge, however, is the efficient training of such AI-based methods for specific systems of interest. In this article, we provide an overview of the training approach being researched at the Professorship of Computer Science in Mechanical Engineering at Helmut-Schmidt University, Hamburg.

**Keywords** – Machine Learning, Molecular Dynamics, Materials Development, Fuel Cells

## ABBREVIATIONS

DFT	Density Functional Theory
MCMC	Monte Carlo Markov Chain

## I. INTRODUCTION

Hydrogen fuel cells are a promising approach for dealing with the growing global demand for clean and sustainable energy. However, the current efficiency and cost of fuel cells pose a significant obstacle to their widespread adoption. The development of new materials to improve fuel cell technology is therefore urgently needed. One example would be the development of a non-precious metal-based catalyst for the oxygen reduction reaction, to replace the currently used platinum-based catalysts, which due to the high cost of platinum make up a significant portion of the total cost of a fuel cell [1].

First-principles atomistic simulations are a valuable tool in the development of novel fuel cell materials. Unfortunately, even though the prediction of chemical and material properties such as catalytic activity from first principles has in principle been possible since the discovery of the equations of quantum mechanics one hundred years ago, it remains a challenging task even to this day due to the computational complexity involved in solving these equations numerically [2]. The simulation of the time evolution of atomistic systems is particularly affected by this computational complexity, since the corresponding equations have to be solved in each timestep of the simulation to evaluate the forces acting on the atoms.

As a consequence, even with the most popular quantum chemical simulation method Density Functional Theory

(DFT) simulations are only possible with a limited time horizon and system size [3]. As a response to these challenges, innovative machine-learning approaches have been developed to address these difficulties. In particular, substitutional modeling of atomistic systems via machine-learning models has become an active area of research [4]. Neural networks, a subset of machine learning, are a promising approach for fast and more scalable simulations while maintaining high-accuracy. This can be achieved by training those neural networks on a dataset consisting of different atomic geometries for which the forces acting on the atoms have been simulated with DFT.

State-of-the-art neural network architectures can now achieve unprecedented accuracy in the prediction of atomic forces and energies with as little as one hundred training geometries on small molecule benchmarks [5] where the training data has been subsampled from molecular dynamics simulations that are long enough, that the molecule traversed its entire configuration space exhaustively. This has been achieved by integrating several important physical properties (e.g. energy conserving forces, symmetry w.r.t. rotations and exchange of identical atoms) directly into the neural network architecture. However, while such benchmarks are useful for evaluating neural network architectures, they often fail to address many of the practical challenges in which computational chemists or materials scientists would employ neural network-based force fields. In particular the problem of how to select atomic arrangements to train the neural network models on.

If classical quantum chemical simulations could already produce molecular dynamics (MD) trajectories of such exhaustive length as the one used to generate the benchmark datasets, there would be no need for a neural network model in the first place. The most impactful applications for neural network models lie in scenarios where classical simulations cannot achieve sufficient time horizons and where the atomic geometries the system might evolve into are unknown in advance.

## II. SOLUTION APPROACH

To enable neural network-accelerated molecular dynamics simulations for fuel cell components (e.g., polyelectrolyte membranes and catalysts), the Professorship of Computer Science in Mechanical Engineering at HSU Hamburg is developing a Bayesian neural network framework based on state-of-the-art neural network architectures. The aim of this framework is to integrate large-scale public materials databases, uncertainty quantification, and classical quantum

chemical simulations into an efficient on-the-fly learning approach.

The approach involves leveraging extensive materials databases containing atomic geometries of thousands of different molecules and materials and their associated force labels to construct a Bayesian prior over the neural network's parameters. This prior is iteratively updated using Bayes' rule as new data from the target system becomes available. During simulations, the Bayesian neural network drives the molecular dynamics until it encounters an atomic arrangement where uncertainty in the prediction exceeds a predefined threshold. At that point, a classical DFT simulation is conducted to compute the forces for the given configuration. These computed forces are then used to update the Bayesian neural network, which resumes driving the molecular dynamics simulation until another high-uncertainty configuration arises (FIGURE 1).

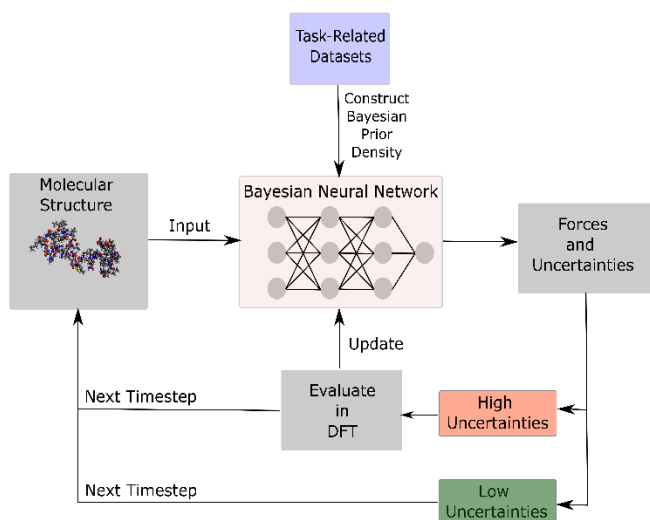


FIGURE 1: ILLUSTRATION OF THE ON-THE-FLY LEARNING APPROACH TO ACCELERATE MOLECULAR SIMULATIONS.

The development of that framework requires three key steps.

First of all, a suitable method for performing Bayesian inference with state-of-the-art neural network models has to be developed and it has to be verified, that it is practically viable (viability of sampling the Bayesian posterior, model accuracy, quality of uncertainty quantification).

As a second step, a suitable Bayesian prior to incorporate large-scale databases has to be identified and empirically verified.

In the last step, the Bayesian neural network model and Bayesian prior have to be integrated into an on-the-fly learning workflow (FIGURE 1).

Currently, the first two steps have been achieved successfully [6,7], while the last step is actively being worked on.

In our paper “High Accuracy Uncertainty-Aware Interatomic Force Modeling with Equivariant Bayesian Neural Networks” we developed a new method for performing Bayesian inference for state-of-the-art neural network architectures used for atomistic modeling [7]. A key challenge with Bayesian models is that they require a process called sampling the posterior distribution, which is a mathematical

way of determining neural network weights that result in a good fit of the training data. Existing methods for this process, such as Monte Carlo Markov Chain (MCMC) sampling, often struggle with modern neural networks because different parts of the network (such as the different layers) have vastly different “scales” of how much the weights influence the predictions of the neural networks. This mismatch can slow down or even prevent the algorithm from converging to a good solution.

To solve this, we developed an improved version of the MCMC algorithm by introducing a parameter-specific step size—a way to adjust how quickly the algorithm learns for each part of the neural network by changing weights that only have a weak influence on the predictions by larger amounts than weights that have a stronger effect. This idea is inspired by techniques already used in traditional machine learning optimizers, and it allowed us to achieve convergence much more efficiently, making Bayesian neural networks practical for real-world use.

After developing this new method, we tested it on several datasets used to benchmark machine learning models in molecular simulations. These datasets included tasks like predicting forces in small molecules and complex polymer chains. We compared our method against other commonly used techniques for estimating uncertainty, such as Monte Carlo Dropout and deep ensembles. Across all tests, our model consistently performed better, delivering more accurate predictions and better estimates of uncertainty.

However, we noticed an interesting issue: all the methods, including ours, tended to be slightly overconfident. This means the uncertainty estimates were sometimes smaller than they should be, leading to occasional underestimation of the potential error in predictions. This behavior, shown in FIGURE 2, highlighted an area for further improvement in future work.

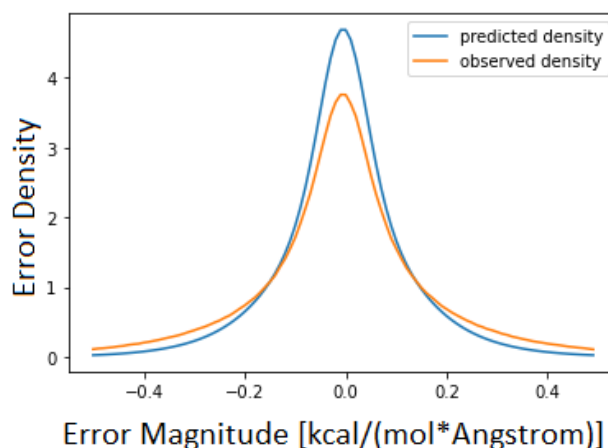


FIGURE 2: OBSERVED AND PREDICTED ERROR DENSITY FOR A TEST DATASET OF AN ASPIRIN MOLECULE.

In our paper, “Bayesian Transfer Learning of Neural Network-Based Interatomic Force Models” we explore how to make the creation of the neural network model smarter and more data efficient [6]. While recent innovations in neural network architectures have made training highly accurate neural network models possible, creating these models can still be challenging because training data still has to be simulated with DFT, which is computationally demanding.

Our goal was to find a way to reduce the number of costly DFT calculations needed to train these models while still maintaining or even improving their accuracy. To achieve this, we utilized transfer learning, where information from one task (or dataset) is reused to help solve another, related task.

Specifically, we focused on how data from two sources could help:

1. Large-scale materials databases, which contain information about many different molecules and materials.
2. Lower accuracy simulations, which are faster but less precise than DFT.

We used this data to create a Bayesian prior, a way of encoding prior knowledge about which neural network weights are likely to work well. This prior acts as a starting point for the model by narrowing down the range of possible weights it considers during training.

Our approach works as follows:

- First, we pre-train the neural network on the large database, allowing it to learn general patterns about interactions between atoms.
- Then, we create a Bayesian prior by giving higher likelihood to neural network weights that are similar to those of the pre-trained model. Weights that are farther away from this starting point are considered less likely to work well.

This method essentially lets the model build on existing or easily producible data reducing the need for extensive new training data from high-accuracy simulations of the specific system being studied.

We tested our approach against previous methods, including one that used no prior knowledge (an "uninformative prior") and a deep ensemble, which is a collection of neural networks trained separately to quantify uncertainty. Our results showed that the model with the constructed prior performed significantly better:

- It achieved higher accuracy with much less training data, making it far more efficient.
- It provided reliable uncertainty estimates, helping identify when its predictions might be less trustworthy.

An example of these results can be seen in FIGURE 3, where the model with the new prior achieves much higher accuracy compared to the other methods.

### III. OUTLOOK

Building on top of the previous two papers, we are currently working on the last remaining challenge, the integration of the developed transfer learning method for fine-tuning pre-trained neural networks into an on-the-fly learning workflow. The key obstacle that has to be overcome for this is the on-the-fly calibration of the uncertainties to overcome the model overconfidence problem.

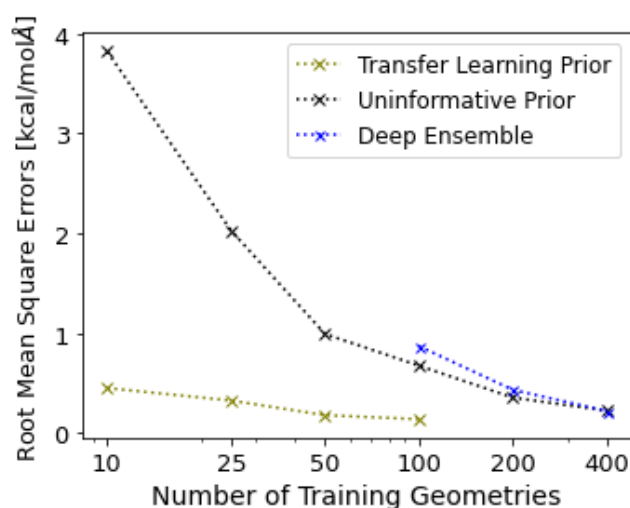


FIGURE 3: COMPARING ACCURACIES OF THE INFORMED AND UNINFORMED PRIOR AS WELL AS A DEEP ENSEMBLE AS A FUNCTION OF THE NUMBER OF GEOMETRIES IN THE TRAINING DATASET.

In a typical machine learning scenario, overconfidence issues can be easily remedied by assessing the degree of overconfidence on a hold-out validation set of data not included in the training data. After assessing the degree of overconfidence on the validation set, the predicted uncertainties on new predictions can be adjusted accordingly, for example by multiplying the standard deviation in the prediction by a fixed factor that was estimated on the validation set.

However, in an on-the-fly learning scenario, there is no validation dataset. The only system specific data available is from DFT calls, when the model uncertainty is too high. A critical issue here is that if at some point in the on-the-fly learning cycle the model becomes too overconfident it might stop calling DFT simulations altogether.

The only additional source of knowledge we have is experience from previous experiments about the degree of overconfidence of the models.

To address these issues, we want to utilize a Bayesian prior over the degree of overconfidence. With this prior we can express our pre-existing knowledge about what the degree of overconfidence typically is. Additionally, it enables us to bias the predicted uncertainties towards underconfidence in the beginning of simulations where only little data is available to accurately estimate the degree of overconfidence for that specific system.

### ACKNOWLEDGMENT

This research as part of the project CouplteIT! is funded by dtec.bw – Digitalization and Technology Research Center of the Bundeswehr which we gratefully acknowledge. dtec.bw is funded by the European Union – NextGenerationEU.

### REFERENCES

- [1] Lixin Fan, Zhengkai Tu, Siew Hwan Chan, "Recent development of hydrogen and fuel cell technologies: A review," *Energy Reports*, Volume 7, p. 8421-8446, 2021.
- [2] Friedman, P.W. Atkins and R.S., "Molecular Quantum Mechanics," OUP Oxford, 2011.
- [3] Marquetand, M. Gastegger and P., "Molecular dynamics with neural network potentials," *Machine Learning Meets Quantum Physics*, p. 233–252, 2020.

- [4] Emir Kocer, Tsz Wai Ko, and Jörg Behler, "Neural network potentials: A concise overview of methods," *Annual Review of Physical Chemistry*, S. 163–186, 2022.
- [5] Simon Batzner, Albert Musaelian, Lixin Sun, Mario Geiger, Jonathan P. Mailoa, Mordechai Kornbluth, Nicola Molinari, Tess E. Smidt, and Boris Kozinsky, "E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials," *Nature Communications*, 2022.
- [6] Tim Rensmeyer, Willi Großmann, Denis Kramer and Oliver Niggemann, "Bayesian Transfer Learning of Neural Network-Based Interatomic Force Models," *The 38th Annual AAAI Conference on Artificial Intelligence | Workshop on AI to Accelerate Science and Engineering*, 2023.
- [7] Tim Rensmeyer, Ben Craig, Denis Kramer, and Oliver Niggemann, "High Accuracy Uncertainty-Aware Interatomic Force Modeling with Equivariant Bayesian Neural Networks," *Digital Discovery*, 2024.