Received: 22 July 2021 | Revised: 28 September 2021 | Accepted: 24 October 2021

DOI: 10.1002/aic.17516

REACTION ENGINEERING, KINETICS AND CATALYSIS

AIChE JOURNAL

Uncertainty quantification in machine learning and nonlinear least squares regression models

Ni Zhan | John R. Kitchin

Department of Chemical Engineering, Carnegie Mellon University, Pittsburgh, PA, USA

Correspondence
John R. Kitchin, Department of Chemical Engineering, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA.
Email: jkitchin@andrew.cmu.edu

Abstract

Machine learning (ML) models are valuable research tools for making accurate predictions. However, ML models often unreliably extrapolate outside their training data. The multiparameter delta method quantifies uncertainty for ML models (and generally for other nonlinear models) with parameters trained by least squares regression. The uncertainty measure requires the gradient of the model prediction and the Hessian of the loss function, both with respect to model parameters. Both the gradient and Hessian can be readily obtained from most ML software frameworks by automatic differentiation. We show examples of the uncertainty method in applications of molecular simulations and neural networks. We further show that the uncertainty measure is larger for input space regions that are not part of the training data. Therefore, this method can be used to identify extrapolation and to aid in selecting training data or assessing model reliability.

KEYWORDS
artificial intelligence, machine learning, uncertainty

1 | INTRODUCTION

Machine learning (ML) models are used in many fields of science and engineering. They can decrease computational time, make predictions and forecasts, and improve insights of complex and high-dimensional datasets. Models are more useful when they provide a prediction with its uncertainty, and in some applications it may be critical to provide a reliable uncertainty estimate.[1,2] The uncertainty estimate should also help identify input data regions that lead to extrapolation. Overall, ML models can be used more reliably when robust uncertainty quantification is available. In the simple case of low-dimensional linear regression, an analytical prediction interval is available,[3] and that can be used to calculate the uncertainty intervals in many statistical software packages. The analytical prediction interval for linear regression requires ($\mathbf{X}^T\mathbf{X}$) to be invertible, where $\mathbf{X} \in \mathbb{R}^{n \times d}$ is the design matrix with n data points and d feature dimensions. In general, for more complex (nonlinear) models and higher dimensional datasets, analytical prediction intervals do not exist and alternative methods are required.

Uncertainty quantification methods for ML models are an active area of research.[4–6] Some common methods are model ensembling and training models with built-in uncertainty such as Gaussian process (GP) regression and quantile regression.[7,8] If we already have a model with trained parameters, we may want to avoid training a different model type or additional model ensemble. In these cases, this article presents a simple uncertainty quantification method for parameterized models trained on minimizing the summed squared error between a model and a dataset. The method exploits automatic differentiation to calculate the Hessian of the loss function based on summed squared errors, and provides an uncertainty estimate which depends on the prediction point, model training data, and the model itself. We show how the uncertainty can identify whether extrapolation is occurring and aids in dataset selection for an example application of molecular simulation. In the remaining sections, we review the background on uncertainty quantification and an application of ML models for molecular simulation, describe the method called the delta method, and show its use in an example neural network (NN) that predicts energies from atomic structures.

1.1 | Uncertainty quantification methods

An ML model is defined by its mathematical definition and training dataset. In complex models, the mathematical definition may be best defined with a computing program (code) and includes the model structure, parameters, and hyperparameters. The training data are most likely preprocessed and modified to normalize and standardize them. The ML models are fitted to the training data, usually by minimizing an error function. In parameterized models, the training process determines the model parameters. We take the model parameters, the error function, exact training dataset with its modifications, and all aspects of the model's mathematical definition to be the ML model.

Here, we describe the motivation and need for the uncertainty estimate. When we use a trained ML model to predict on another dataset, the prediction accuracy depends on the training data used. If the new data are in an extrapolation region, the accuracy is not expected to be good.[9] It is not obvious when the model is predicting in an extrapolation region if the ML model has a high-dimensional input space. A motivation for an uncertainty measurement is to help determine when the model is extrapolating, or is simply not reliable. Knowledge of the uncertainty helps in identifying when a model is reliable and indicates confidence in a model prediction.

Uncertainty quantification methods for models include bootstrap, ensembling, using model-specific uncertainty, and the delta method. Bootstrap, ensembling, and the delta method can be used for parametric models and NNs. Bootstrap uncertainty is based on statistical theory, has some different variations, and requires training multiple models on different bootstrap samples of the data or residuals.[10–12] With the different models, the uncertainty on predictions can be estimated.

Ensemble methods require training multiple models on the entire dataset, and the different models give the uncertainty.[13–15] The uncertainty estimate quantitatively improves as the ensemble size increases, so the optimal number of ensemble models is unknown and must be user-determined.[16] Some studies used bootstrap uncertainty for NN models and showed that it gave more reliable uncertainties than the delta method,[17,18] but training multiple models is computationally expensive and time consuming.

Model-specific uncertainty include GP regression,[19] dropout for NNs,[20] and Bayesian NNs.[21–23] Obtaining uncertainties in this way limits the possible mathematical forms of the model. Dropout and Bayesian NNs are also more difficult to train than standard NNs. The next section describes specific instances in which ensembling, bootstrap, and GP methods were used for molecular simulation. In another method, the posterior of model parameters can be approximated by the Laplace approximation, which is the second-order Taylor series expansion around the optimal model parameters.[24,25] Hence, the Hessian of the log-likelihood contains relevant information about uncertainty.

This article focuses on the delta method, a method which also uses the Hessian of the log-likelihood. The method is based on linearly approximating the model and uses an estimate of the standard error of model parameters assuming maximum likelihood estimation. Therefore, the method applies to models with parameters trained by minimizing squared error, and the model structure could be a simple linear regression to complex nonlinear regression including NNs. Different variants of the method use approximations of the Hessian, such as the outer product of gradients, and experiments have tested the delta method on NN models.[17,26–30] We further describe the method in Methods, with theoretical details in Appendix S1.

1.2 | Addressing uncertainty in molecular simulation

We examine uncertainty quantification using the example application of molecular simulation, an area which has benefited from ML. Here, we describe the background of ML for molecular simulation. Simulations allow researchers to obtain materials' physical properties and quickly screen materials. Molecular dynamics (MD) and Monte Carlo simulations require a model of the potential energy surface (PES) which predicts energies and forces from atomic configurations. The options for the PES model include first-principles methods such as density functional theory (DFT), physical potentials, and ML potentials. ML potentials aim to achieve the high accuracy of DFT at a significantly faster computation. ML potentials are also more systematically improvable than physical potentials.[31] Many studies have successfully used ML potentials in simulations.[32–35]

Uncertainty quantification is useful for ML potentials. Commonly used ML potentials such as NNs will usually unreliably extrapolate on inputs much different from their training data. A consequence of extrapolation during a molecular simulation is that it likely gives wrong or unphysical results. The best ways to select enough of the relevant training space are nonobvious, since the space of atomic structures is often large, not well understood, and not possible to enumerate. Furthermore, atomic structures are translated into fingerprints which are high dimensional and less human-interpretable than the original atomic configurations. Hence, we require a method to determine the uncertainty of a prediction from an ML potential, and the quantitative uncertainty helps us to avoid extrapolation and identify sparse regions in the training dataset.

Current methods developed to address uncertainty are ensembles of potentials, on-the-fly methods, and using ML models with built-in uncertainty. Ensemble methods independently train two or more ML potentials, and check for agreement between them. In Behler's approach, NNs with different architectures are trained, and atomic structures whose predictions' differ significantly across NNs are added to the training set.[36–38] Peterson et al. trained an ensemble of 50 NN potentials and found that ensemble spread was a good indicator for prediction error across the space.[39] Smith et al. also used ensemble disagreement to approximate prediction error and select a small training set.[40,41]

In MD simulations, on-the-fly methods use an ML potential augmented with quantum mechanical (QM) calculations.[42] There is a query if the ML prediction can be used. If it fails, a QM calculation is run and added to a database, and the ML model can be retrained. A simple query is whether the fingerprint is out of the minimum and maximum bounds in the current database.[43] This is a minimum requirement that the ML model is not extrapolating; however, guaranteeing that the fingerprint is within bounds of training data does not guarantee a low error.[44]

Another approach is training GP regressions[42,45] or other ML models with built-in uncertainty estimates. Vandermause et al. and Xie et al. used GP uncertainty to train potentials on the fly.[46,47] Many ML models, such as NNs, do not have theoretical guarantees for uncertainty of a prediction. Perturbation of NN weights could provide some range of uncertainty.[48] Another work used dropout in NN training as a Bayesian approximation and thereby calculating uncertainties for interatomic potentials.[49] Janet et al. used the distance in values of the last layer of NNs (or latent space) as an uncertainty measure.[50] Tran et al. compared GP, Bayesian NN, dropout NN, and ensembles of different NN structures and found more conservative uncertainties for GP and overconfident uncertainties for Bayesian NN, dropout, and NN ensemble.[51]

Musil et al. compared GP, ensembling with random subsets of the data, and bootstrap methods for obtaining uncertainties of predicting formation energies on molecular datasets.[52] They found that random sampling was easier to implement than bootstrapping, computationally faster than GP for uncertainty estimates, and matches the true error and uncertainty from GP. Li et al. trained NN potentials with different NN structures (number of nodes), weight initialization, and learning rates, and compared the resulting prediction accuracies.[53] Their work showed a quantitative uncertainty arising from some NN hyperparameters, but it required training several NN potentials for a new system, and does not provide confidence or prediction intervals. In an alternative approach, Botu et al. fitted an empirical function to an uncertainty estimate as a function of fingerprint distance between input and reference training fingerprints.[54] Their uncertainty estimation approach requires a larger training set size.

Overall, there is no clear consensus on the best uncertainty quantification method, and its selection usually depends on the model form used, for example, built-in uncertainty from GPR or ensembles when using NNs. The delta method provides a simple alternative for providing quantitative uncertainty when a pretrained model exists, without the necessity of training additional models. That is the focus of this article.

2 | METHODS

The delta method applies to regression problems of a model g with parameters θ. The residuals of model prediction are assumed to be Gaussian-distributed and centered around zero. We assume that the model parameters $\widehat{\theta}$ were obtained by minimizing a function of the summed squared errors, although the method can be extended to maximize *a posteriori* estimation and cross-entropy loss for classification tasks.[55] We obtain an approximate standard error of a model prediction $g\left(\widehat{\theta},x\right)$ by using a Taylor series approximation and an approximate standard error of $\widehat{\theta}$. Suppose that $\frac{\partial g\left(\widehat{\theta},x,\right)}{\partial\widehat{\theta}}$ is nonzero, then the standard error of $g\left(\widehat{\theta},x\right)$ for a point x is given in Equation (1) using the delta method.

$$\mathrm{se}\left(g\left(\widehat{\theta},x\right)\right)\approx\sqrt{\frac{\partial g\left(\widehat{\theta},x\right)}{\partial\widehat{\theta}}^{T}I_n^{-1}\frac{\partial g\left(\widehat{\theta},x\right)}{\partial\widehat{\theta}}} \tag{1}$$

where $\frac{\partial g\left(\widehat{\theta},x,\right)}{\partial\widehat{\theta}}$ is the gradient vector of the model with respect to parameters at the point x for which we are calculating uncertainty, and I_n is the Fisher information matrix, defined as the expectation of the Hessian of the negative log-likelihood. The Fisher information is related to the Hessian of the loss, usually the sum of squared errors, by a scaling factor. Equation (1) shows that the model prediction of the standard error is a function of the training data, model, and point for which the uncertainty is calculated.

For small to medium models, the delta method is faster and easier to implement compared to ensembling, and the Hessian and gradients of the model are readily obtained with automatic differentiation that is included in most ML packages. To demonstrate the ease of use, we show a simple code example using the `autograd`[56] package in Listing 1 (results in Figure 1). The delta method is limited by model size because the Hessian will be $m \times m$, where m is number of parameters, and the Hessian needs to be inverted. For very large models and NNs, inverse Fisher information approximations have been proposed. Ritter et al. implemented Kronecker-factored Hessian for a network with around 2 million parameters,[24] and Nilsen et al. proposed and implemented eigenvalue spectrum Hessian approximation on networks with around 100,000 parameters.[55] For even larger networks, other uncertainty quantification methods would likely be needed.

The uncertainties are calculated after the model has finished training, and the Fisher information inverse needs to be calculated only once per model and training dataset. In previous tests, a model with 861 parameters and 1900 training data points required around 5 min to calculate the inverse Fisher information with an Intel Core i7-7820HQ CPU @ 2.9 GHz using `autograd`. Using more modern automatic differentiation, frameworks are expected to be faster. Calculating the uncertainties after obtaining the inverse Fisher information requires much less time. Theoretically, calculation of the Fisher information matrix scales quadratically with number of parameters and linearly with number of training data points.

The quality of standard errors calculated using the delta method depends on the fit of the model. We found that well-fitted models have better uncertainty measures, and our assumptions required residuals to be independent and identically distributed normal around zero. Poorly fitted models have uncertainty measures that are less quantitatively accurate.

2.1 | Practical modifications to the inverse Fisher matrix

There are a few steps or best practices to modify how the inverse Fisher information matrix is computed.

1. We start with H, the Hessian of the sum squared errors loss function. For some models, such as NNs, the Hessians of the loss functions with respect to parameters are often nearly singular with some eigenvalues much larger than the others,[57,58] and the optimal parameters may be at a saddle point.
2. Add a small number ϵ to the Hessian diagonal. Adding ϵ to the diagonals makes the matrix better conditioned for taking its inverse. ϵ should be larger in magnitude than the most negative eigenvalue. We used $\epsilon = \max(1e-5, 1.05 \cdot \text{abs}(\lambda_{\min}(H)))$, where $\lambda_{\min}(H)$ is the smallest eigenvalue of H. Modifying the Hessian of the objective function with respect to NN parameters has been suggested in the literature and is justified because the top eigenvalues are a few orders of magnitude larger than the other eigenvalues.[24,58,59] Also note that Hessian conditioning suggests that the number of parameters of the NN is much larger than the actual degrees of freedom of the NN.
3. We take the Moore–Penrose pseudoinverse $(H + \epsilon I)^{-1}$. If the inverse exists, which is most cases following step 2, it is the same as the true inverse.
4. Multiply $(H + \epsilon I)^{-1}$ by a scaling factor α. This is done to calibrate the uncertainties to be near the residuals. We set α to be mean squared error (MSE) in most cases. To select α, we suggest trying n^{β}·MSE, where n is the number of training data points and β is any nonnegative number, but usually in the range [0, 1]. α should be chosen so that uncertainties have the same order of magnitude as the residual errors for the training dataset.
5. Force the scaled inverse $P := \alpha(H + \epsilon I)^{-1}$ to be positive semi-definite. For eigendecomposition $P = Q\Lambda Q^{-1}$, the closest positive semi-definite matrix in terms of Frobenius norm is $Q\max(\Lambda, 0)Q^{-1}$, where max is element-wise max.[60]

The final inverse Fisher information I_n^{-1} used in Equation (1) is $Q\max(\Lambda, 0)Q^{-1}$.

2.2 | Code example

We show a simple code implementation of the delta method in Listing 1 using the `autograd` package.[56] Note that obtaining the required gradients is a single line for the Hessian (line 27) and gradients (line 34), demonstrating the ease of automatic differentiation. Similar codes would apply in `PyTorch`[61] or other ML packages. In this simple example, we fit a quadratic function to some slightly noisy data and show the resulting confidence intervals on the fit (Figure 1). The Hessian in this case was well-conditioned, so the modifications described above were not necessary.

LISTING 1 Autograd example of the delta method

```
1  import autograd.numpy as np
2  from autograd import elementwise_grad, hessian
3  from scipy.optimize import minimize
4  import matplotlib.pyplot as plt
5  from scipy.stats.distributions import t
6
7  # {x, y} data
8  x = np.array([0.1, 0.3, 0.5, 0.7, 0.9])
9  y = np.array([0.0, 0.1, 0.3, 0.5, 0.8])
10
11 def g(theta, x):
12     'function with parameters theta'
13     return theta[0] * x**2 + theta[1] * x + theta[2]
14
15 def sse(theta):
16     'Summed squared error objective function'
17     return np.sum((g(theta, x) - y)**2)
18
19 initial_guess = np.array([0.1, 0.5, 0.2])
20 # minimize sse
21 sol = minimize(sse, initial_guess)
22 theta = sol.x
23 ypred = g(theta, x)
24
25 # obtain Hessian of sse using autograd
26 h = hessian(sse)(theta)
27 # inverse and scale Hessian
28 p = sse(theta) / len(x) * np.linalg.pinv(h)
29
30 uncerts = []
31 for xi in x:
32     # obtain gradient using autograd
33     gprime = elementwise_grad(g, 0)(theta, xi)
34     # delta method
35     uncerts += [np.sqrt(gprime @ p @ gprime)]
36 uncerts = np.array(uncerts)
37
38 # t-value
39 tval = t.ppf(0.975, len(x))
40
41 # plot the data, fit, and confidence intervals
42 plt.figure(figsize=(3.25, 4))
43 plt.plot(x, y, 'o')
44 plt.plot(x, ypred)
45 plt.plot(x, ypred + tval * uncerts, '--r')
46 plt.plot(x, ypred - tval * uncerts, '--r')
47 plt.xlabel('x')
48 plt.ylabel('y')
49 plt.legend(['Data', 'Prediction', '95% confidence']
50 plt.tight_layout()
51 plt.savefig('simple-code-ex.png')
```

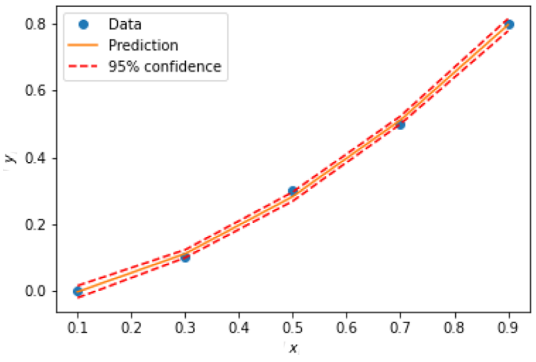

FIGURE 1 Result from Listing 1

This simple example shows all the pieces of the delta method. There are data and a function (line 11) with parameters that are fitted to the data. The regression here is done by optimization (line 21); this problem is linear and could be solved analytically, but we show the optimization approach for generality. We used automatic differentiation to obtain the Hessian (line 26) and gradient of the function (line 33) with respect to the parameters. The rest is conventional linear algebra.

In calculating the *t*-value (line 39), technically the degrees of freedom should be used instead of the number of data points. However for large NNs, the effective degrees of freedom is much smaller than the number of model parameters, and determining NN degrees of freedom is an active area of research.[62–64] For simplicity, we used the number of data points to estimate the *t*-value throughout our results.

3 | RESULTS

We show examples of using the delta method on different models to demonstrate how the uncertainty behaves. We begin with a simple one-dimensional (1-D) NN, and build in complexity in subsequent examples.

3.1 | One dimension input NN

This example is a 1-D input NN. We start with 1-D input for clearer intuition and visualization. We generated synthetic data from the one-dimensional Lennard–Jones (LJ) function and added some Gaussian noise. We fitted these data to an NN with structure [1, 4, 1] (one input, one hidden layer with four nodes, and one output) using `scipy.optimize.minimize`. The NN had 13 total parameters.

We test how the standard error changes with different training datasets. We generated two sets of training data to fit the NN, and Figure 2 shows the fits. These sets of training data were from the same LJ function and had the same variance of Gaussian noise added. We expect the true function to be within the confidence interval 95% of the time. In Figure 2A, the uncertainty increases for large and small *x*, which is desirable because we do not know how the NN will behave in those regions outside the training data. In Figure 2B, there is a region of missing data in the middle, and the confidence interval expands in the region of missing data. These cases demonstrate that the uncertainty depends on the training data in a useful way. The uncertainty generally increases in regions with less data, which makes sense because we are less certain of our model in a space with less training data.

3.2 | High-dimensional NN potential

3.2.1 | Trained NN potential

This example applies the delta uncertainty method to a high-dimensional NN potential. We use the SingleNN (implemented in `PyTorch`) and weighted Behler–Parinello style symmetry functions.[65–67] The data are DFT energy and force calculations based on atomic configurations, specifically the dataset used in Boes 2017,[68] which contains 3907 unique AuPd slabs. The symmetry functions transform the atomic configuration information into a vector of numbers, or "fingerprint," and we used four weighted G_2 symmetry functions. For the NN, we used two hidden layers with 11 nodes each; thus the NN architecture is [4, 11, 11, 1], which is 211 total parameters.

To demonstrate the usefulness of the uncertainty method, we start by training on a subset of the data. This mimics the iterative approach often used in training these models. We then check for extrapolation on the remaining data using the delta method. For this first potential, 572 configurations with a lattice constant of 3.934 Å were randomly split into 64%, 16%, and 20% training, validation, and test sets, respectively. The NN was trained on energies and atomic forces using SingleNN, and uncertainties were calculated in the same `PyTorch` framework.

Figure 3 shows the energy parity plots of the training, validation, and test sets. The parity is good in all cases, and root mean squared

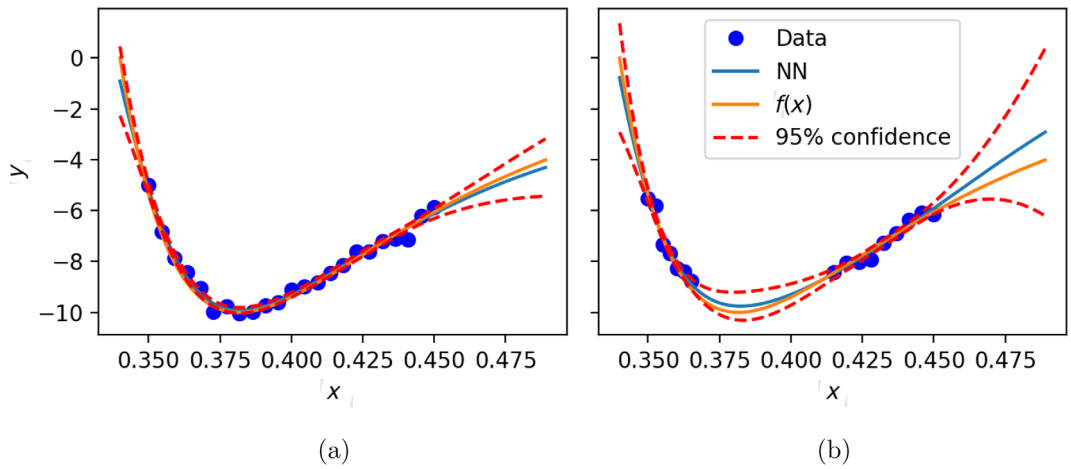

FIGURE 2 One dimension input NN and confidence intervals. (A) Twenty-three training data points, and confidence interval is wider at the edges. (B) Region of missing data in middle, and confidence interval expands in region of missing data

errors (RMSEs) are 0.003, 0.0023, and 0.003 eV/atom for train, validation, and test, respectively. Figure 4 shows the distributions of standard errors of confidence for the three datasets. These distributions are very similar and mostly overlapping. Figure 5 shows the parity plot of the test set with 95% prediction intervals. The true values are within the prediction intervals for 98% of the dataset, which is close to 95% and shows that the delta method provides quantitatively reasonable uncertainties in this case. The main result is that similar datasets with the same accuracy using the model will have similar distributions of uncertainties.

Next, we use the same potential to predict on a new dataset. If the new dataset is dissimilar from the training data, we expect the uncertainties to be high. While all the training set had 3.934 Å lattice constants, the new dataset has 4.034 and 4.134 Å lattice constants, which we will refer to as predict-4.0 and 4.1 datasets. As a result, we expect the fingerprints to differ from those of the train set, that is, we

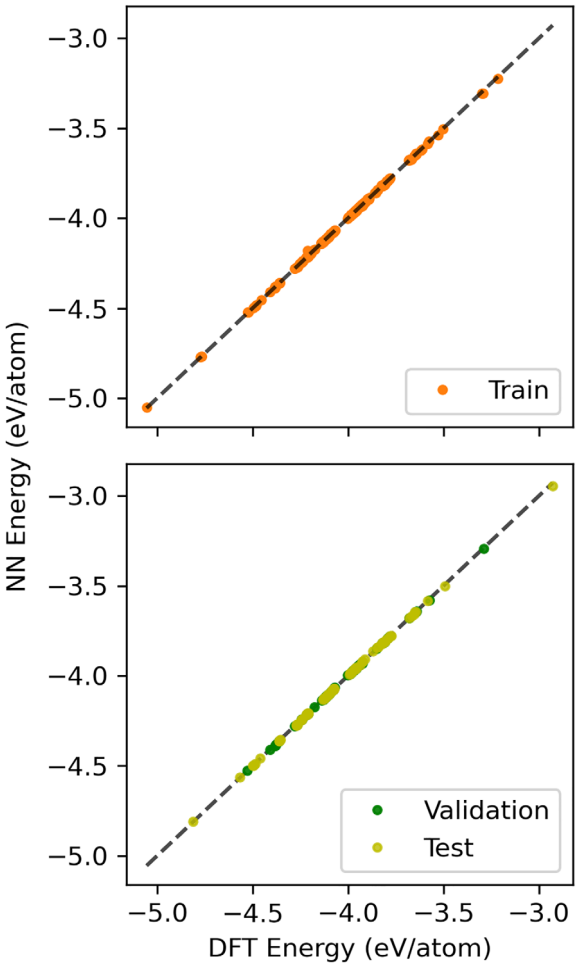

FIGURE 3 Parity plot of SingleNN

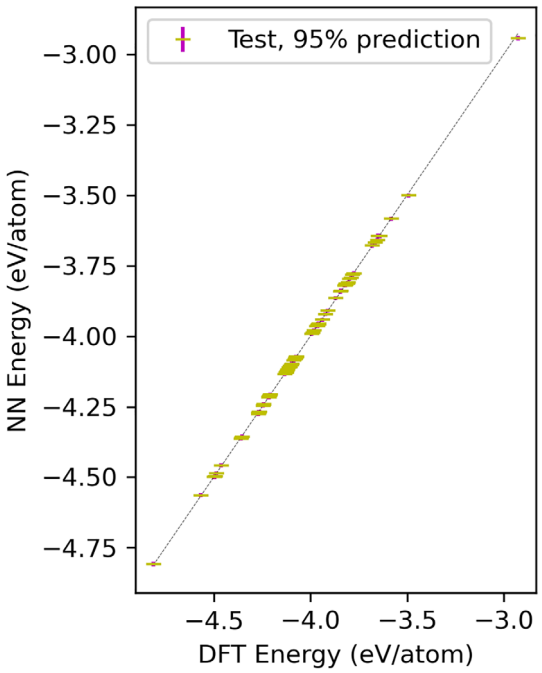

FIGURE 5 Parity plot with 95% prediction intervals for test set

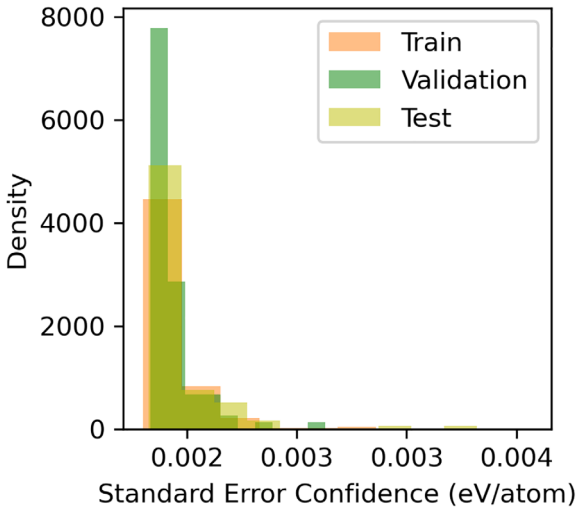

FIGURE 4 Distribution of uncertainties (standard error of confidence)

know we are extrapolating here. Figure 6 shows the energy parity plots for the predict sets with 95% prediction intervals. The predictions are offset with an error, and the uncertainties are clearly much larger than those for the test set from Figure 5. Table 1 shows the average standard error of confidence/prediction for the datasets. Training and test sets have around the same standard error confidence of 0.002 eV/atom, and predict-4.0 and 4.1 sets have higher uncertainties of 0.023 and 0.034 eV/atom, respectively, which are one order of magnitude larger than that of training and test. Since this uncertainty is much larger, it could indicate that the model is extrapolating on the predict sets, and the parity plots (Figure 6) seem to indicate this.

We examine the fingerprints, and Figure 7 shows an example fingerprint for the train and predict datasets. There are regions where the predict-4.0 and 4.1 atoms' fingerprints are outside the training distributions, which is suggestive of extrapolation. For predict-4.0, the true values are within the prediction intervals for 75% of the dataset, which is not that close to 95%; however for predict-4.1, the true values are within the prediction intervals for 0% of the dataset. This

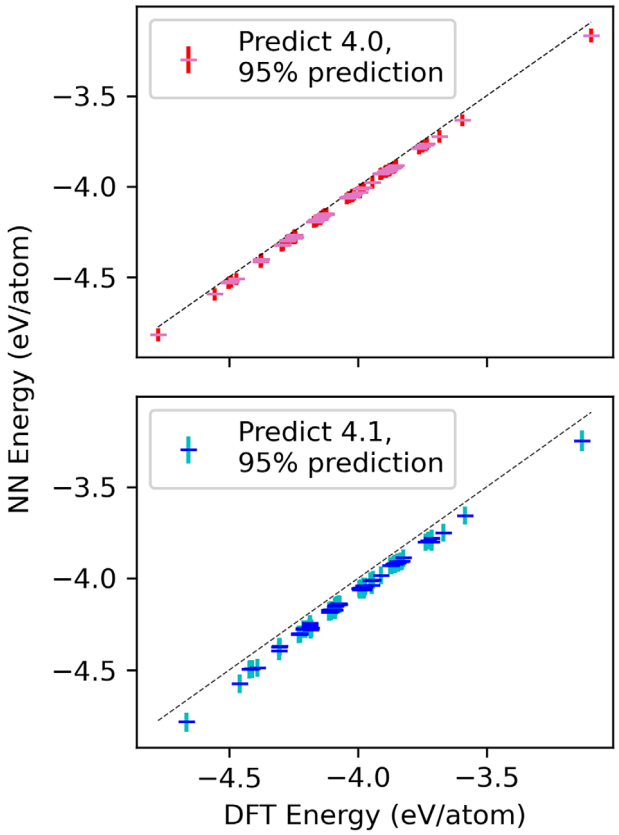

FIGURE 6 Prediction on new lattice datasets; uncertainty may be much larger in an extrapolation region

TABLE 1 Average standard error confidence of datasets

Dataset	Average standard error confidence (eV/atom)	Average standard error prediction (eV/atom)
Test	0.0020	0.0036
Predict-4.0	0.0234	0.0235
Predict-4.1	0.0336	0.0337

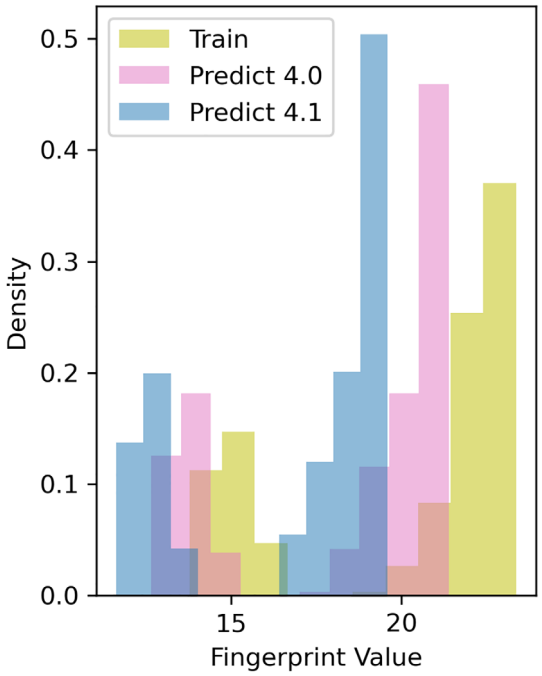

FIGURE 7 The predict-4.0 and 4.1 datasets have fingerprints outside the range of training distribution (fingerprint example shown is eta = 0 with Pd center atoms)

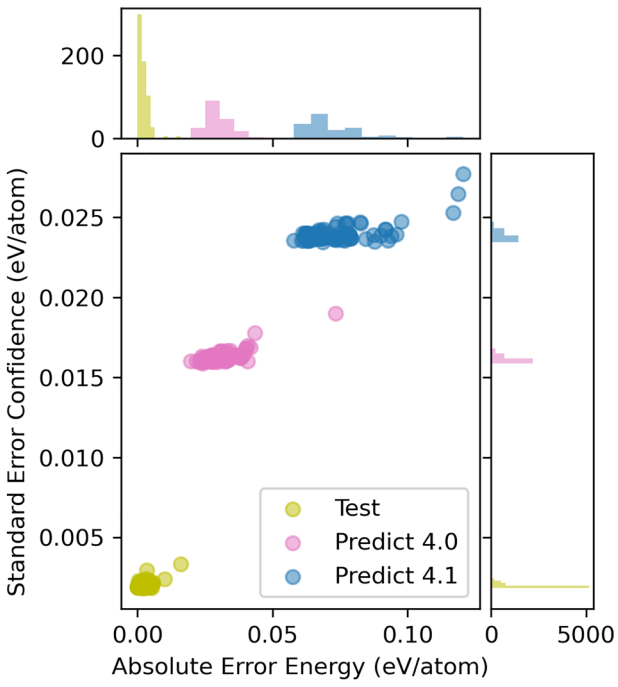

FIGURE 8 Standard error from delta method vs. absolute error and their distributions

seems to indicate that the prediction interval becomes less quantitatively accurate as the extrapolation increases. However, when the uncertainty is much larger than the training uncertainties, the model is likely extrapolating, and we should not trust the prediction. Therefore, this uncertainty method helps in identifying the data regions where a model extrapolates. Figure 8 shows the standard error confidence vs. absolute energy error, and their distributions for the test and predict datasets. Figure 8 shows the general trend that uncertainty from the delta method increases when the true error increases. The trend is most obvious in a heterogeneous dataset.

3.2.2 | Uncertainties after retraining

Next, we retrain the potential with some of the predict-4.0 and 4.1 data and check how uncertainties are affected. We expect the uncertainties to decrease after retraining. We added 64% of each predict-4.0 and 4.1 dataset, or 365 data points each, and retrained. Figure 9 shows the energy parity plots of the new training and predict sets. After retraining, the predict set is on parity and no longer offset. The true values are within the prediction intervals for 98.7% of the training data and 98.5% of the predict data, which are close to the theoretical 95% and show that the uncertainties calculated from the delta method are quantitatively reasonable. Figure 10 shows the updated standard error confidence vs. absolute energy error and their distributions for test and predict datasets. After retraining, the standard errors across datasets are mostly overlapping, and the average standard errors are the same for the datasets. The average standard error confidence and predict are 0.002 and 0.003 eV/atom, respectively. Since we retrained on the predict-4.0 and 4.1 datasets, we are no longer extrapolating on that data and the uncertainties updated to reflect this: they are no longer an order of magnitude larger than the train sets' as was the case before retraining. We can use this uncertainty method to iteratively retrain a potential by adding data with high uncertainties. This is sometimes called active learning.

In the calculation of the Fisher information matrix, we used the errors of energies only, although we trained on energies and forces. From a theoretical perspective, the Fisher information should include some information about force errors, but exactly how much to include is not obvious. By using only loss of energies, we save computational time for calculating the Fisher information, and the uncertainty measurement still accomplishes the objective and is quantitatively reasonable. Therefore in practice, using only the loss of energies for the Fisher information works well.

We can also extend uncertainty to other properties such as forces. For this case, in Equation (1), $g\left(\widehat{\theta}\right)$ is force, which is $-\frac{\partial E}{\partial \text{position}}$, where E represents energy. We obtain $g'\left(\widehat{\theta}\right)$ through automatic differentiation by taking the derivative of $-\frac{\partial E}{\partial \text{position}}$ with respect to model parameters. In this way, we use the delta method to calculate uncertainties

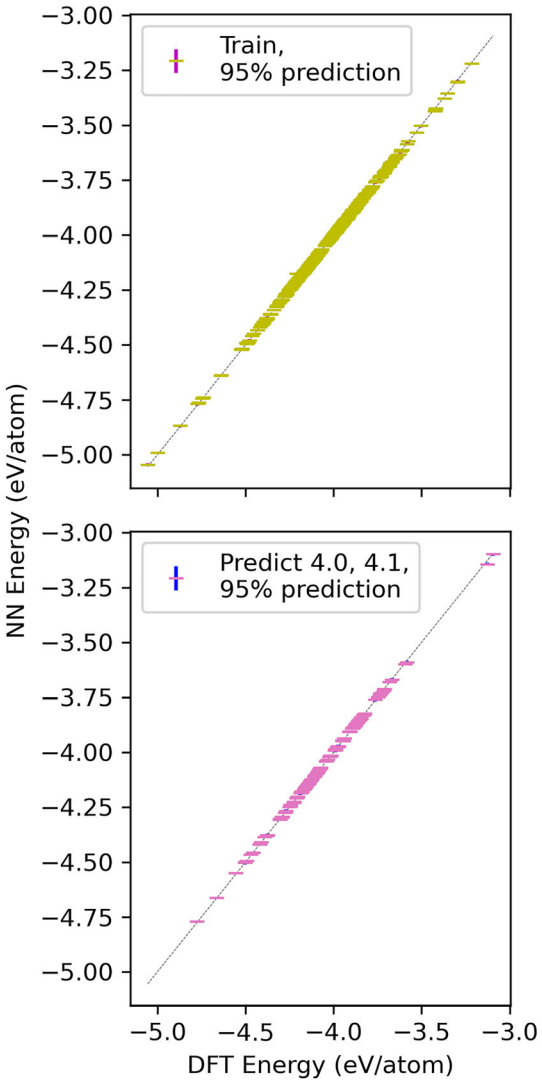

FIGURE 9 Parity plot after retraining

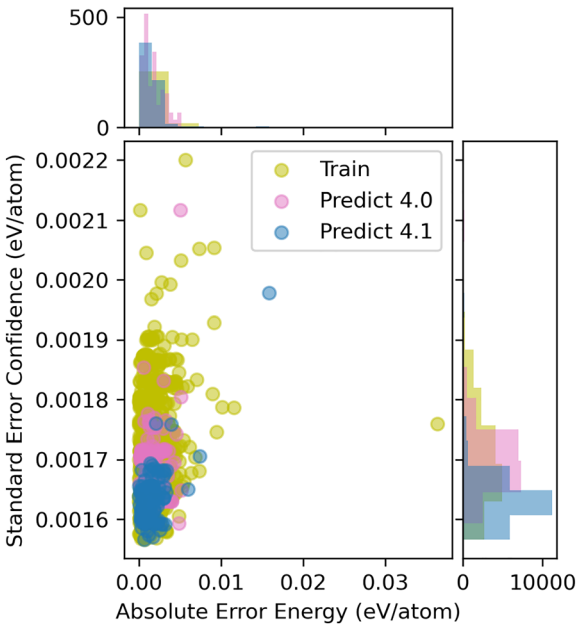

FIGURE 10 Distribution of uncertainties after retraining

for other quantities of interest. Further work can be done to investigate the quality and methods for force uncertainties of NN potentials.

There is a possibility for fast approximations of the Fisher information after retraining. If we retrain by adding one or a few new training points, we may want a cheaper calculation of the Fisher information matrix. The Fisher information matrix can be linearly separated by training data since the loss is a sum over training data points. If the parameters of the model did not change from retraining, then the new Fisher information is the summation of the original Fisher information and the Fisher information for the new training points. Since retraining likely alters the model parameters, the previous Fisher information from old model parameters is an approximation. If only a few training points are added and the model parameters do not change much, taking the Fisher information of the new training points and adding it to the original can be a fast approximation of the true Fisher information. Further work is required to determine when this approximation is adequate.

4 | CONCLUSIONS

The delta method is a fast and easy way to estimate uncertainty. It requires the Hessian of the loss and gradient with respect to model parameters, and these are obtainable with most ML packages using automatic differentiation. The delta method is applicable to most models that are parametric and have nonzero gradients with respect to parameters. The uncertainty estimate will depend on the training data, model, and input (point) for which the uncertainty is calculated. The delta method is an alternative to ensemble or bootstrapping methods for obtaining uncertainty estimates, and uncertainty estimates are important because they can help determine when a model is extrapolating and increase model reliability.

We showed an application of the delta method to a high-dimensional NN potential in molecular simulation. We illustrated how we can iteratively retrain a model by adding data with high uncertainties to improve it. This could also be done on the fly, for example, while running an MD simulation with an ML potential. The uncertainty can determine the longest timescale MD simulation that is valid for a potential, or identify when additional data should be added to the training data to improve it. The utility of the delta method shown here extends far beyond molecular simulation; it can also be applied to many other applications that rely on linear or nonlinear regression models.

AUTHOR CONTRIBUTIONS

Ni Zhan: Methodology (lead); writing – original draft (lead). **John R. Kitchin:** Conceptualization (equal); project administration (equal); writing – review and editing (equal).

DATA AVAILABILITY STATEMENT

The data that supports the findings of this study are available in the supplementary material of this article

ORCID

John R. Kitchin https://orcid.org/0000-0003-2625-9232

REFERENCES

1. Ovadia Y, Fertig E, Ren J, et al. Can you trust your model's uncertainty? Evaluating predictive uncertainty under dataset shift. CoRR; 2019. http://arxiv.org/abs/1906.02530v2
2. Kompa B, Snoek J, Beam AL. Second opinion needed: communicating uncertainty in medical machine learning. *NPJ Digit Med*. 2021;4(1):4. doi:10.1038/s41746-020-00367-3
3. Wasserman L. *All of Statistics*. Springer Texts in Statistics. Springer; 2004. doi:10.1007/978-0-387-21736-9
4. Maddox W, Garipov T, Izmailov P, Vetrov D, Wilson AG. A simple baseline for Bayesian uncertainty in deep learning. CoRR; 2019. http://arxiv.org/abs/1902.02476v2
5. Izmailov P, Maddox WJ, Kirichenko P, Garipov T, Vetrov D, Wilson AG. Subspace inference for Bayesian deep learning. In: Adams RP, Gogate V, eds. *Proceedings of the 35th Uncertainty in Artificial Intelligence Conference, Vol. 115 of Proceedings of Machine Learning Research*. PMLR; 2020:1169-1179. http://proceedings.mlr.press/v115/izmailov20a.html
6. Liu JZ, Lin Z, Padhy S, Tran D, Bedrax-Weiss T, Lakshminarayanan B. Simple and principled uncertainty estimation with deterministic deep learning via distance awareness. CoRR; 2020. http://arxiv.org/abs/2006.10108v2
7. Lei J, G'Sell M, Rinaldo A, Tibshirani RJ, Wasserman L. Distribution-free predictive inference for regression. CoRR; 2016. http://arxiv.org/abs/1604.04173v2
8. Romano Y, Patterson E, Candès EJ. Conformalized quantile regression. CoRR; 2019. http://arxiv.org/abs/1905.03222v1
9. Hajinazar S, Shao J, Kolmogorov AN. Stratified construction of neural network based interatomic models for multicomponent materials. *Phys Rev B*. 2017;95(1):014114. doi:10.1103/physrevb.95.014114
10. Endo T, Watanabe T, Yamamoto A. Confidence interval estimation by bootstrap method for uncertainty quantification using random sampling method. *J Nucl Sci Technol*. 2015;52(7–8):993-999. doi:10.1080/00223131.2015.1034216
11. Palmer G, Du S, Politowicz A, et al. Calibrated bootstrap for uncertainty quantification in regression models. CoRR; 2021. http://arxiv.org/abs/2105.13303v1
12. Du H, Barut E, Jin F. Uncertainty quantification in CNN through the bootstrap of convex neural networks. *Proc AAAI Conf Artif Intell*. 2021;35(13):12078-12085. https://ojs.aaai.org/index.php/AAAI/article/view/17434
13. Abdar M, Pourpanah F, Hussain S, et al. A review of uncertainty quantification in deep learning: techniques, applications and challenges. *Inf Fusion*. 2021;76:243-297. doi:10.1016/j.inffus.2021.05.008
14. Wenzel F, Snoek J, Tran D, Jenatton R. Hyperparameter ensembles for robustness and uncertainty quantification. CoRR; 2020. http://arxiv.org/abs/2006.13570v3
15. Lakshminarayanan B, Pritzel A, Blundell C. Simple and scalable predictive uncertainty estimation using deep ensembles. CoRR; 2016. http://arxiv.org/abs/1612.01474v3
16. Berger JO, Smith LA. On the statistical formalism of uncertainty quantification. *Ann Rev Stat Appl*. 2019;6(1):433-460. doi:10.1146/annurev-statistics-030718-105232
17. Tibshirani R. A comparison of some error estimates for neural network models. *Neural Comput*. 1996;8(1):152-163. doi:10.1162/neco.1996.8.1.152
18. Dybowski R, Roberts SJ. Confidence intervals and prediction intervals for feed-forward neural networks. *Clinical Applications of Artificial Neural Networks*. Cambridge University Press; 2001:298-326.
19. Rasmussen C, Williams C. *Gaussian Processes for Machine Learning*. MIT Press 2006. www.GaussianProcess.org/gpml

20. Gal Y, Ghahramani Z. Dropout as a Bayesian approximation: representing model uncertainty in deep learning. CoRR; 2015. http://arxiv.org/abs/1506.02142v6
21. Mackay DJC. *Bayesian methods for adaptive models*. PhD thesis. California Institute of Technology; 1992.
22. Neal RM. Priors for infinite networks. *Bayesian Learning for Neural Networks*. Springer; 1996:29-53. doi:10.1007/978-1-4612-0745-0_2
23. Blundell C, Cornebise J, Kavukcuoglu K, Wierstra D. Weight uncertainty in neural networks. CoRR; 2015. http://arxiv.org/abs/1505.05424v2
24. Ritter H, Botev A, Barber D. A scalable Laplace approximation for neural networks. Paper presented at: 6th International Conference on Learning Representations, ICLR 2018-Conference Track Proceedings, Vol. 6. International Conference on Representation Learning; 2018.
25. Tian H, Rzepa C, Upadhyay R, Rangarajan S. Estimating vibrational and thermodynamic properties of adsorbates with uncertainty using data driven surrogates. *AIChE J*. 2019;65(12). doi:10.1002/aic.16838e16838.
26. Donaldson JR, Schnabel RB. Computational experience with confidence regions and confidence intervals for nonlinear least squares. *Dent Tech*. 1987;29(1):67-82. doi:10.1080/00401706.1987.10488184
27. de Veaux RD, Schumi J, Schweinsberg J, Ungar LH. Prediction intervals for neural networks via nonlinear regression. *Dent Tech*. 1998;40(4):273. doi:10.2307/1270528
28. Papadopoulos G, Edwards P, Murray A. Confidence estimation methods for neural networks: a practical comparison. *IEEE Trans Neural Netw*. 2001;12(6):1278-1287. doi:10.1109/72.963764
29. Settles B. Active learning literature survey; 2009.
30. Khosravi A, Nahavandi S, Creighton D, Atiya AF. Comprehensive review of neural network-based prediction intervals and new advances. *IEEE Trans Neural Netw*. 2011;22(9):1341-1356. doi:10.1109/tnn.2011.2162110
31. Kitchin JR. Machine learning in catalysis. *Nat Catal*. 2018;1(4):230-232. doi:10.1038/s41929-018-0056-y
32. Boes JR, Kitchin JR. Neural network predictions of oxygen interactions on a dynamic Pd surface. *Mol Simul*. 2017;43(5–6):346-354. doi:10.1080/08927022.2016.1274984
33. Shakouri K, Behler J, Meyer J, Kroes GJ. Accurate neural network description of surface phonons in reactive gas-surface dynamics: N_2 + Ru(0001). *J Phys Chem Lett*. 2017;8(10):2131-2136. doi:10.1021/acs.jpclett.7b00784
34. Kobayashi R, Giofré D, Junge T, Ceriotti M, Curtin WA. Neural network potential for Al-Mg-Si alloys. *Phys Rev Mater*. 2017;1(5):053604. doi:10.1103/physrevmaterials.1.053604
35. Deringer VL, Bernstein N, Bartók AP, et al. Realistic atomistic structure of amorphous silicon from machine-learning-driven molecular dynamics. *J Phys Chem Lett*. 2018;9(11):2879-2885. doi:10.1021/acs.jpclett.8b00902
36. Behler J. Representing potential energy surfaces by high-dimensional neural network potentials. *J Phys Condens Matter*. 2014;26(18):183001. doi:10.1088/0953-8984/26/18/183001
37. Behler J. Constructing high-dimensional neural network potentials: a tutorial review. *Int J Quantum Chem*. 2015;115(16):1032-1050. doi:10.1002/qua.24890
38. Behler J. First principles neural network potentials for reactive simulations of large molecular and condensed systems. *Angew Chem Int Ed*. 2017;56(42):12828-12840. doi:10.1002/anie.201703114
39. Peterson AA, Christensen R, Khorshidi A. Addressing uncertainty in atomistic machine learning. *Phys Chem Chem Phys*. 2017;19(18):10978-10985. doi:10.1039/c7cp00375g
40. Smith JS, Nebgen B, Lubbers N, Isayev O, Roitberg AE. Less is more: sampling chemical space with active learning. *J Chem Phys*. 2018;148(24):241733. doi:10.1063/1.5023802
41. Smith JS, Nebgen BT, Zubatyuk R, et al. Approaching coupled cluster accuracy with a general-purpose neural network potential through transfer learning; 2019. https://chemrxiv.org/articles/Outsmarting_Quantum_Chemistry_Through_Transfer_Learning/6744440
42. Li Z, Kermode JR, Vita AD. Molecular dynamics with on-the-fly machine learning of quantum-mechanical forces. *Phys Rev Lett*. 2015;114(9):096405. doi:10.1103/physrevlett.114.096405
43. Botu V, Ramprasad R. Adaptive machine learning framework to accelerate ab initio molecular dynamics. *Int J Quantum Chem*. 2014;115(16):1074-1083. doi:10.1002/qua.24836
44. Artrith N, Behler J. High-dimensional neural network potentials for metal surfaces: a prototype study for copper. *Phys Rev B*. 2012;85(4):045439. doi:10.1103/physrevb.85.045439
45. Bartók AP, Kermode J, Bernstein N, Csányi G. Machine learning a general-purpose interatomic potential for silicon. *Phys Rev X*. 2018;8(4):041048. doi:10.1103/physrevx.8.041048
46. Vandermause J, Torrisi SB, Batzner S, et al. On-the-fly active learning of interpretable Bayesian force fields for atomistic rare events. *NPJ Comput Mater*. 2020;6(1):20. doi:10.1038/s41524-020-0283-z
47. Xie Y, Vandermause J, Sun L, Cepellotti A, Kozinsky B. Bayesian force fields from active learning for simulation of inter-dimensional transformation of stanene. *NPJ Comput Mater*. 2021;7(1):40. doi:10.1038/s41524-021-00510-y
48. Christensen R. *Error mitigation in computational design of sustainable energy materials*. PhD thesis. Department of Energy Conversion and Storage, Technical University of Denmark; 2016.
49. Wen M, Tadmor EB. Uncertainty quantification in molecular simulations with dropout neural network potentials. *NPJ Comput Mater*. 2020;6(1):124. doi:10.1038/s41524-020-00390-8
50. Janet JP, Duan C, Yang T, Nandy A, Kulik H. A quantitative uncertainty metric controls error in neural network-driven chemical discovery. *Chem Sci*. 2019;10:7913-7922. doi:10.1039/c9sc02298h
51. Tran K, Neiswanger W, Yoon J, Zhang Q, Xing E, Ulissi ZW. Methods for comparing uncertainty quantifications for material property predictions. *Mach Learn Sci Technol*. 2020. doi:10.1088/2632-2153/ab7e1a1(2):025006.
52. Musil F, Willatt MJ, Langovoy MA, Ceriotti M. Fast and accurate uncertainty estimation in chemical machine learning. *J Chem Theory Comput*. 2019;15:906-915. doi:10.1021/acs.jctc.8b00959
53. Li Y, Xiao W, Wang P. Uncertainty quantification of artificial neural network based machine learning potentials. *Volume 12: Materials: Genetics to Structures*; 2018. 10.1115/imece2018-88071
54. Botu V, Batra R, Chapman J, Ramprasad R. Machine learning force fields: construction, validation, and outlook. *J Phys Chem C*. 2016;121(1):511-522. doi:10.1021/acs.jpcc.6b10908
55. Nilsen GK, Munthe-Kaas AZ, Skaug HJ, Brun M. Epistemic uncertainty quantification in deep learning classification by the delta method. CoRR; 2019. http://arxiv.org/abs/1912.00832v2
56. autograd: https://github.com/HIPS/autograd
57. Sagun L, Bottou L, LeCun Y. Eigenvalues of the hessian in deep learning: singularity and beyond CoRR; 2016. http://arxiv.org/abs/1611.07476v2
58. Gur-Ari G, Roberts DA, Dyer E. Gradient descent happens in a tiny subspace. CoRR; 2018. http://arxiv.org/abs/1812.04754v1
59. Gill J, King G. What to do when your hessian is not invertible. *Sociol Methods Res*. 2004;33(1):54-87. doi:10.1177/0049124103262681
60. Cheng SH, Higham NJ. A modified Cholesky algorithm based on a symmetric indefinite factorization. *SIAM J Matrix Anal Appl*. 1998;19:1097-1110.
61. PyTorch. https://pytorch.org/
62. Larsen BW, Fort S, Becker N, Ganguli S. How many degrees of freedom do we need to train deep networks: a loss landscape perspective. CoRR; 2021. http://arxiv.org/abs/2107.05802v1

63. Ingrassia S, Morlini I. Equivalent number of degrees of freedom for neural networks. In: Decker R, Lenz HJ, eds. *Advances in Data Analysis*. Springer; 2007:229-236.
64. Gao T, Jojic V. Degrees of freedom in deep neural networks. *Proceedings of the Thirty-Second Conference on Uncertainty in Artificial Intelligence, UAI'16*. AUAI Press; 2016:232-241. https://www.auai.org/uai2016/proceedings/papers/257.pdf
65. Behler J. Atom-centered symmetry functions for constructing high-dimensional neural network potentials. *J Chem Phys*. 2011;134(7):074106. doi:10.1063/1.3553717
66. Gastegger M, Schwiedrzik L, Bittermann M, Berzsenyi F, Marquetand P. wACSF-weighted atom-centered symmetry functions as descriptors in machine learning potentials. *J Chem Phys*. 2018;148(24):241709. doi:10.1063/1.5019667
67. Liu M, Kitchin JR. Singlenn: modified Behler-Parrinello neural network with shared weights for atomistic simulations with transferability. *J Phys Chem C*. 2020;124(32):17811-17818. doi:10.1021/acs.jpcc.0c04225
68. Boes JR, Kitchin JR. Modeling segregation on AuPd(111) surfaces with density functional theory and Monte Carlo simulations. *J Phys Chem C*. 2017;121(6):3479-3487. doi:10.1021/acs.jpcc.6b12752

SUPPORTING INFORMATION

Additional supporting information may be found in the online version of the article at the publisher's website.

How to cite this article: Zhan N, Kitchin JR. Uncertainty quantification in machine learning and nonlinear least squares regression models. *AIChE J*. 2021;e17516. doi:10.1002/aic.17516