

PAPER • OPEN ACCESS

## Clarifying trust of materials property predictions using neural networks with distribution-specific uncertainty quantification

To cite this article: Cameron J Guich *et al* 2023 *Mach. Learn.: Sci. Technol.* **4** 025019

View the [article online](#) for updates and enhancements.

You may also like

- [Comprehensive Evaluation of Electricity Market Based on Analytic Hierarchy Process and Evidential Reasoning Methods](#)  
Z X Jing, J H Shi, Z Y Luo et al.
- [Model of Optimization solution Based on Generalized Evidence Theory](#)  
Wu Ying
- [Robust and scalable uncertainty estimation with conformal prediction for machine-learned interatomic potentials](#)  
Yuge Hu, Joseph Musielewicz, Zachary W Ulissi et al.



## PAPER

## OPEN ACCESS

## Clarifying trust of materials property predictions using neural networks with distribution-specific uncertainty quantification

RECEIVED  
7 February 2023REVISED  
23 March 2023ACCEPTED FOR PUBLICATION  
5 April 2023PUBLISHED  
3 May 2023Cameron J Gruich<sup>1,2</sup> , Varun Madhavan<sup>1</sup> , Yixin Wang<sup>3</sup>  and Bryan R Goldsmith<sup>1,2,\*</sup><sup>1</sup> Department of Chemical Engineering, University of Michigan, Ann Arbor, MI 48109-2136, United States of America<sup>2</sup> Catalysis Science and Technology Institute, University of Michigan, Ann Arbor, MI 48109-2136, United States of America<sup>3</sup> Department of Statistics, University of Michigan, 1085 S University Ave, Ann Arbor, MI 48109-1107, United States of America

\* Author to whom any correspondence should be addressed.

E-mail: [bgoldsm@umich.edu](mailto:bgoldsm@umich.edu)**Keywords:** computational catalysis, crystal graph convolutional neural networks, evidential regression, recalibration, calibrationSupplementary material for this article is available [online](#)Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

**Abstract**


It is critical that machine learning (ML) model predictions be trustworthy for high-throughput catalyst discovery approaches. Uncertainty quantification (UQ) methods allow estimation of the trustworthiness of an ML model, but these methods have not been well explored in the field of heterogeneous catalysis. Herein, we investigate different UQ methods applied to a crystal graph convolutional neural network to predict adsorption energies of molecules on alloys from the Open Catalyst 2020 dataset, the largest existing heterogeneous catalyst dataset. We apply three UQ methods to the adsorption energy predictions, namely  $k$ -fold ensembling, Monte Carlo dropout, and evidential regression. The effectiveness of each UQ method is assessed based on accuracy, sharpness, dispersion, calibration, and tightness. Evidential regression is demonstrated to be a powerful approach for rapidly obtaining tunable, competitively trustworthy UQ estimates for heterogeneous catalysis applications when using neural networks. Recalibration of model uncertainties is shown to be essential in practical screening applications of catalysts using uncertainties.

**Abbreviations**

ML	machine learning
UQ	uncertainty quantification
CGCNN	crystal graph convolutional neural network
DER	deep evidential regression
OCP	Open Catalyst Project
OC20	Open Catalyst 2020 Dataset
NN	neural network
DFT	density functional theory
MAE	mean absolute error
RMSE	root mean square error
MDAE	median absolute error
MARPD	mean absolute relative percent difference
$C_v$	coefficient of variation
MC	Monte Carlo
GPR	Gaussian process regression

**1. Introduction**

ML approaches have rapidly grown in popularity to accelerate catalyst screening and understanding [1–4]. Making predictions of catalyst properties with ML models is orders of magnitude faster compared to first-principles simulation of catalysts, for example, using DFT modeling to compute adsorption energies of



molecules. DFT modeling combined with ML has emerged as a compelling approach for rapid materials characterization, enabling a several-orders-of-magnitude expansion in the number of materials able to be studied compared to DFT modeling alone [5–7]. However, the potential of ML models to efficiently explore large catalyst spaces can only be achieved if it is simple to detect when ML model predictions are accurate or highly uncertain.

Uncertainty provides a means to infer the accuracy of a model without explicitly knowing the accuracy, but often this uncertainty of the ML model cannot be trusted [8]. Reliable UQ of ML model predictions is a fundamental challenge to ML-guided materials and molecule discovery [8–12]. Predictive uncertainty can be estimated with either distribution-specific methods or distribution-free methods (e.g. MC dropout [13] with a Gaussian-specific calibration assumption [14] or conformal UQ [15] respectively). By distribution-specific methods, we mean methods that assume an underlying distribution shape to the uncertainty of model predictions or their calibrated scaling, while distribution-free methods make no such assumption. Established UQ methods for ML models are often costly to obtain and have limitations in assessing prediction errors for chemical space exploration [16].

Effective uncertainty estimates of predictions are an important design element to enhance several advanced ML strategies for materials discovery, such as high-throughput screening [17, 18], transfer learning [19], and active learning [20, 21], as illustrated in figure 1. For catalyst discovery, one can ascertain if an ML model is making accurate predictions of catalyst behavior by confirming the predicted properties experimentally or through first-principles calculations. In a high-throughput context where many thousands of systems or more are being studied, frequently confirming the predictive accuracy of the ML model via experiment or first-principles calculations is impractical. UQ is a means by which one can quantify the trustworthiness of ML predictions in a way that is practical and attainable for materials discovery strategies that explore vast materials spaces.

ML-guided studies of catalysts that apply uncertainty have already been performed, but these studies historically rely on GPR [22–25]. GPR is formulated from Bayesian statistics and directly outputs uncertainty estimates of the predictions [26, 27]. However, the computational cost of training a GPR model typically scales  $O(N^3)$  and thus grows unfavorably with dataset size, which is an open challenge for big-data applications [28]. Therefore, there is a practical limitation of these models for high-throughput discovery strategies.

NNs are popular because of their relatively high accuracy across diverse chemical spaces and excellent performance for large datasets [29–33]. For the OC20 dataset developed under the open catalyst project (OCP) [34], state-of-the-art deep NN model variants (e.g. CGCNN [35], SchNet [36], and DimeNet++ [37]) have shown up to a 40% improvement in their accuracy of adsorption energy predictions by using 460 000 training samples versus 10 000 samples [34]. The arrival of large training datasets—and innovations in NN model architectures—has led to steady improvements in accuracy for catalyst property predictions,

such as molecular adsorption energies [38]. However, less attention in the field has been devoted to understanding and quantifying the predictive uncertainty associated with these NN architectures applied to prevailing catalysis datasets [14, 15].

For uncertainty-guided sample acquisition and efficient model training for catalysis using NNs, it is crucial to ensure that the uncertainty estimates from a given UQ method are reliable and useful. Typical metrics for evaluating uncertainty estimates include accuracy, sharpness, dispersion, calibration, and tightness, which are discussed in section 2.4. The reliability of a UQ estimate is often primarily discussed in terms of calibration, which refers to how well the uncertainty estimate represents a confidence interval that encloses the ground-truth target [39–44]. Calibrated uncertainty estimates can be used to forecast the accuracy of materials predictions, thereby quantifying the trustworthiness of an ML model to examine and propose materials for further study. Using calibrated uncertainty estimates to forecast whether a large set of ML predictions are reliably accurate or not is common to many applications, such as autonomous driving [45], image analysis [46], and materials discovery [47].

Herein, we compare three different UQ methods, namely,  $k$ -fold ensembling [48–50], MC dropout [13], and evidential regression [51, 52] to better ensure effective use of NN model architectures for high-throughput catalyst discovery strategies. All three methods are able to express uncertainty as a standard deviation ( $\sigma$ ). Evidential regression has seen use in small molecule prediction [51] but has not yet been explored for catalysis applications to our knowledge. These three UQ methods are studied using a CGCNN to predict adsorption energies of molecules on solid catalyst materials based on the OC20 dataset, many of which are binary and ternary alloys [34]. Alloy catalysts are highly relevant in industrial applications, such as the Haber–Bosch process [53], catalytic cracking of hydrocarbons [54], and naphtha reforming [55]. The OC20 dataset is an ideal test case for understanding UQ methods because it is the largest and most diverse heterogeneous catalysis dataset, which represents a case study for many material search challenges.

We use accuracy, sharpness, dispersion, calibration, and tightness as UQ metrics to compare the trustworthiness of each of the three UQ methods, figure 2. We define a trustworthy UQ method as one that is perfectly calibrated because calibration measures how well the uncertainty intervals probabilistically grow or shrink with the predictive error. In practice, one may not know if a UQ method is perfectly calibrated; moreover, it is possible for two UQ methods to have the same average measure of calibration but have different uncertainty estimates for the same materials, so accuracy, sharpness, dispersion, and tightness are discussed as secondary trustworthiness criteria when comparing UQ methods. We particularly emphasize differences in calibration between UQ methods via adversarial group calibration [56] to determine their appropriateness for high-throughput catalyst discovery strategies. Scalar recalibration [57] is applied to address poor calibration performance between UQ methods. Evidential regression is found to be competitively trustworthy before recalibration and the most trustworthy after recalibration. We demonstrate the use of UQ and evidential regression to enumerate materials for DFT-predicted adsorption in the case of a hydrogen adsorbate. Improved reliability in the uncertainty estimates for this demonstration is observed after recalibration. This work will guide future efforts of uncertainty-guided catalyst search and discovery using ML.

## 2. Experimental methods

### 2.1. Dataset: OC20

We used the OC20 dataset to train and test our CGCNN model to predict adsorption energies of molecules on catalyst surfaces [34]. OC20 encompasses a state-of-the-art collection of DFT calculations of adsorption energies on binary and ternary alloy catalysts spanning the periodic table, as well as pure metals. The dataset is comprised of 82 nitrogen, oxygen, and carbon-containing adsorbates. OC20 has over 130 000 000 data points associated with 5243 unique alloy catalyst compositions.

Within OC20, there are different versions of datasets depending on what task an ML model is performing to predict adsorption energy. We performed the initial-structure-to-relaxed-energy (IS2RE) task, where the ML model accepts an unrelaxed initial structure of an adsorbate/alloy system and predicts the relaxed energy of the system. Electronic energies of the geometry-optimized (i.e. relaxed) gas phase species and the bare alloy surface are subtracted from the combined adsorbate/alloy system electronic energy to obtain an adsorption energy at 0 K [34]. We trained the CGCNN models on the entire IS2RE training dataset composed of 460 328 systems. To benchmark the model accuracy and the effectiveness of the UQ methods, test adsorption energy predictions were made on the entire IS2RE Validation In-Domain (Val-ID) dataset composed of 24 943 systems. The Val-ID dataset was composed of adsorbate/alloy systems that the model has never seen before during training. We herein refer to the Val-ID dataset as the test dataset, and we use this dataset as the test dataset because the dedicated test dataset provided within OC20 does not include target labels due to its usage in leaderboard competitions.

## 2.2. Model: CGCNN

We employ the CGCNN model architecture provided by OC20 for adsorption energy predictions [34]. The CGCNN in our study is a deep-learning convolutional NN of atomic surface structures using the atoms and distances between atoms in the material crystal structure as node and edge information, respectively [35]. The node and edge information are combined to form a crystal graph, a graph-based representation of the underlying material. Gaussian basis functions are also used to encode the distances between atoms on the crystal graphs [34]. Element-specific properties of the materials standard to OC20 were employed as node embeddings [34]. This crystal graph and associated properties passes through convolutional layers that reduce the incoming high-dimensional space to a fixed lower-dimensional space. The model accepts adsorbate/alloy systems of a varying number of atoms because all the data in transit inside the model gets convolved to the same number of dimensions. After the convolutional layers, the data was transformed by a series of fully connected layers. Model parameters were optimized to produce accurate predictions via mini-batch stochastic gradient descent. The best reported hyperparameters reported in the original OC20 dataset report were used for our CGCNN model [34]. Before applying any UQ method, our baseline CGCNN model had a 0.651 eV MAE on the IS2RE Val-ID dataset, which is comparable to the expected error reported in the literature [34].

## 2.3. UQ methods

We compared three different UQ methods from the literature:  $k$ -fold ensembling [48–50], MC dropout [13], and evidential regression [51, 52], figure 2(a). Each adsorption energy prediction was expressed as a mean prediction  $\mu$  with uncertainty as a standard deviation  $\sigma$  centered around the prediction. Uncertainty intervals around each prediction were constructed using the associated value of the standard deviation  $\sigma$  to give an interval of  $\mu \pm 3\sigma$ . From an application standpoint, trustworthy uncertainty intervals are interpreted similarly to a confidence interval in that these intervals quantify the lack of confidence in ML predictions. Unlike confidence intervals, uncertainty intervals only have coverage guarantees if a UQ method is perfectly calibrated, figure 2(c).

### 2.3.1. $k$ -fold ensembling

The  $k$ -fold ensemble method estimates uncertainty by segregating the training dataset into  $k$  nonoverlapping subsets (i.e. folds) that are used to train  $k$  models. Each model makes adsorption energy predictions on the same test set, but the predictions for these same systems will be different from model to model because each model was trained on a different fold of training data. For every adsorbate/alloy system,  $k$  predictions were averaged together to produce a mean adsorption energy prediction  $\mu$ . The uncertainty (i.e. a standard deviation  $\sigma$ ) was calculated for each average prediction  $\mu$ . Herein, we used  $k = 5$ .

### 2.3.2. MC dropout


MC dropout estimates uncertainty by modifying the fully connected layers of the NN model. Dropout was applied both during and after training when making predictions—this protocol is necessary to approximate a Gaussian process [13]. Like the ensemble method, all the adsorption energy predictions corresponding to any specific system were used to calculate a mean adsorption energy prediction  $\mu$  and an associated uncertainty  $\sigma$  for that system.

We refer to the different adsorption energy predictions for the same adsorbate/alloy system as MC dropout samples. When we say an uncertainty estimate is a 1000-sample MC dropout estimate, we mean that 1000 different adsorption energy predictions for the same system were used post-training to calculate a mean prediction  $\mu$  and uncertainty estimate  $\sigma$ . For each fully connected layer of our CGCNN model, 5% of the nodes were stochastically dropped out for each MC dropout sample. A dropout rate of 5% was chosen because it results in a negligible reduction in model accuracy on the test set, figure S1.

### 2.3.3. Evidential regression

Evidential regression estimates uncertainty by modifying both the loss function and model architecture of a regression model [51, 52]. Evidential regression does not require sampling predictions post-training to estimate the uncertainty, unlike  $k$ -fold ensembling and MC dropout. All the information needed to make an uncertainty estimate of a prediction is outputted by the model at prediction time. Evidential regression uses an evidential loss function  $L_i(w)$  such that the model learns the predictive uncertainty while it minimizes the predictive error during training

$$L_i(w) = L_i^{\text{NLL}}(w) + \lambda L_i^{\text{R}}(w). \quad (1)$$



Here  $w$  refers to the model weights that are optimized during training,  $L_i^{\text{NLL}}(w)$  refers to the negative log-likelihood (NLL) of the evidential loss function, and  $L_i^R(w)$  refers to the regularization term of the evidential loss function. These terms are defined as follows:

$$L_i^{\text{NLL}}(w) = \frac{1}{2} \log\left(\frac{\pi}{v}\right) - \alpha \log(\Omega) + \left(\alpha + \frac{1}{2}\right) \log\left((y_i - \gamma)^2 v + \Omega\right) + \log\left(\frac{\Gamma(\alpha)}{\Gamma(\alpha + \frac{1}{2})}\right) \quad (2)$$

$$L_i^R(w) = |y_i - \gamma| \cdot (2v + \alpha). \quad (3)$$

Here  $\Omega = 2\beta(1 + v)$  and the parameters  $\gamma$ ,  $v$ ,  $\alpha$ , and  $\beta$  are the evidential distribution parameters. The symbol  $\Gamma$  refers to the gamma function. The variable  $\lambda$  is a scalar hyperparameter that controls the degree of regularization introduced by the term  $L_i^R(w)$ . We modified the CGCNN model architecture to predict these four parameters for each adsorbate/alloy system. The parameter  $\gamma$  represents the predicted target property (here the adsorption energy of a system). The parameters  $v$ ,  $\alpha$ , and  $\beta$  are related to the aleatoric ( $\sigma_a$ ) and epistemic ( $\sigma_e$ ) uncertainties of each prediction:

$$\sigma_a = \frac{\beta}{\alpha-1} \quad (4)$$

$$\sigma_e = \frac{\beta}{v(\alpha-1)}. \quad (5)$$

Whereas  $k$ -fold ensembling and MC dropout provide one uncertainty estimate for each prediction, evidential regression provides two estimates corresponding to either aleatoric uncertainty or epistemic uncertainty. For this experiment, we focus on the epistemic uncertainty of evidential regression because epistemic uncertainty is related to the uncertainty of the ML model and is a reducible uncertainty, whereas aleatoric uncertainty is a type of stochastic, irreducible uncertainty [58]. Both aleatoric and epistemic uncertainty estimates across values of  $\lambda$  are included in figures S2 and S3 of the SI for completeness.

The regularization term  $L_i^R(w)$  inflates the uncertainty estimate of the adsorption energy predictions based on the absolute residual error  $|y_i - \gamma|$ , and  $\lambda$  weights the magnitude of this inflation. Because  $\lambda$  was fixed before training the model and inflates the uncertainty of the predictions, evidential regression gives a tunable uncertainty estimate that grows or shrinks depending on the choice of  $\lambda$ .

## 2.4. UQ metrics

We compare the UQ methods in terms of five UQ metrics: accuracy, sharpness, dispersion, calibration, and tightness [14, 44]. These UQ metrics collectively provide a means of comparing the predictive accuracy, shape, and reliability of the uncertainty distribution between UQ methods. Ideally, a UQ method would be accurate, sharp, disperse, calibrated, and tight, figure 2(b). A UQ method that expresses a prediction as a mean and the uncertainty as a standard deviation should give accurate predictions regardless of the uncertainty centered around the predictions. A UQ method is ideally sharp, implying that the method gives highly certain adsorption energy predictions on average. A UQ method is preferably disperse such that the uncertainty level of different test predictions is easily distinguished. Importantly, the UQ method should also be calibrated to allow reliable interpretation as a confidence interval. Lastly, a UQ method should be tight in that the predictive uncertainty interval should only be as large as necessary to capture the ground truth.

### 2.4.1. Accuracy

The MAE, RMSE, MDAE, MARPD, coefficient of determination ( $R^2$ ), and the correlation coefficient ( $R$ ) were used to quantitatively assess the accuracy of each UQ method. Because MDAE uses the median of the absolute residual error distribution across predictions, it is less sensitive to outliers compared to the RMSE. MAE and MARPD are formulated as means, so these metrics have some sensitivity to outliers. MARPD is a normalized measure of accuracy [14]:

$$\text{MARPd} = \frac{1}{N} \sum_{i=1}^N 100 * \frac{|\hat{y}_i - y_i|}{|\hat{y}_i| + |y_i|} \quad (6)$$

where  $\hat{y}_i$  is the  $i$ th adsorption energy prediction from the UQ method applied onto our CGCNN model and  $y_i$  is the true adsorption energy of the  $i$ th adsorbate/alloy system as predicted by DFT.

### 2.4.2. Sharpness

Sharpness (Sha) was expressed as [14]:

$$\text{Sha} = \sqrt{\frac{1}{N} \sum_{i=1}^N \sigma_i^2} \quad (7)$$

where  $N$  is the total number of data points and  $\sigma_i^2$  is the squared uncertainty of the  $i$ th adsorption energy prediction.

### 2.4.3. Dispersion

The dispersion of uncertainty across predictions for each UQ method was measured by constructing a box plot for each method and measuring the interquartile range (IQR)

$$\text{IQR} = Q3 - Q1. \quad (8)$$

$Q3$  is the third quartile corresponding up to the 75th percentile and  $Q1$  is the first quartile corresponding up to the 25th percentile. Whiskers are calculated based on  $1.5 \times \text{IQR}$  below the first quartile and above the third quartile, respectively. Each box plot was overlaid with a violin plot, which provides a smooth kernel density estimate of the distribution shape such that each distribution can be visually inspected similarly to a

histogram. Each kernel density estimate was performed using Scott's rule for calculating the estimator bandwidth [59].

$C_V$  was calculated to measure dispersion as well. While IQR quantifies the spread of a distribution bulk on an absolute scale,  $C_V$  quantifies the spread of a distribution relative to its mean. These metrics in tandem allow for a more robust analysis of dispersion than either alone.  $C_V$  was expressed as:

$$C_V = \frac{\sigma_\sigma}{\mu_\sigma} \quad (9)$$

where  $\sigma_\sigma$  and  $\mu_\sigma$  respectively are the standard deviation and mean of the uncertainty distribution. We include Bessel's correction in the calculation of  $\sigma_\sigma$ .

#### 2.4.4. Calibration

Calibration formally refers to how well the uncertainty estimate represents the true correctness likelihood of a prediction [39, 41–44]. A calibrated ML model is a model whose uncertainty estimates should be comparable with its predictive error residuals [14]. In other words, a calibrated model should often be highly certain for highly accurate predictions and vice versa. Quantitatively, a calibrated model means that the uncertainty magnitude  $\sigma$  of an adsorption energy prediction  $\mu$  should often be comparable in magnitude to the residual predictive error ( $y - \mu$ ), where  $y$  is the ground-truth adsorption energy of the system the model tries to correctly predict.

Calibration of a model was assessed by constructing a calibration curve, figure 2(b). A calibration curve displays the true frequency of points in each confidence interval relative to the predicted fraction of points in that interval [60]. For calibration curve construction, a quantile-based method was used [56]. Programmatically, this task was accomplished in the same way as explained by Tran *et al* [14].

The residual predictive error was normalized by the uncertainty of that prediction. All the normalized residual errors of all the predictions were combined to construct a distribution. If this distribution was comparable to a unit Gaussian distribution, then the uncertainty across all the predictions was said to be calibrated on average. In other words, the uncertainty across all predictions is calibrated on average if the normalized residual predictive error follows the probability density function of a unit Gaussian distribution:

$$\Phi \left( z = \frac{y - \mu}{\sigma} \right) = \frac{e^{-\frac{z^2}{2}}}{\sqrt{2\pi}} \quad (10)$$

where  $z$  is the normalized residual predictive error. This distribution-specific measure of calibration was applied to all UQ methods tested.

To better interpret this procedure with intuition, consider the case where an adsorption energy prediction is highly accurate but has high uncertainty. In this case, the normalized residual error ( $z$ ) is close to zero because the numerator is small; likewise, the denominator (i.e. uncertainty) is sufficiently large. It is desirable to have a normalized residual less than or equal to unity for predictions because this implicitly shows that the uncertainty interval is large enough to enclose the ground truth adsorption energy. The deviation of the normalized residual error from unit Gaussian behavior is what produces a nonideal calibration curve, which is constructed by defining quantiles and comparing the proportion of data inside each quantile between distributions.

Every calibration curve shown herein visually represents the average calibration, which refers to calibration across the entire adsorption energy test set. However, randomly selected subgroups of predictions and even individual predictions should be calibrated as well [56, 60]. We refer to the global scope of calibration across all possible subgroups as the calibration density; to the best of our knowledge, we have not seen this term used in the literature. We use this term to better intuitively emphasize that any stochastic subselection of adsorption energy predictions across varying scales should be calibrated, that is, any subselection of predictions should ideally produce a normalized residual error distribution that is unit Gaussian, figure S4. If this expectation is met, then we say that a UQ method has effective calibration density. Calibration density is a related concept to individual uncertainty estimates in that it conceptually represents the map of calibration at all sample sizes and quantifies the deviation in the average uncertainty from individual estimates [61].

Calibration density was assessed via adversarial group calibration [56, 61]. Ten subgroups of predictions were drawn across different subgroup sizes (e.g. ten subgroups that each represent 10% of the test set, 20%, 30%, etc). Within each subgroup, predictions were drawn without replacement. For each subgroup size, the ten subgroups were compared in an adversarial fashion. By adversarial, we mean that the worst performing subgroup of the ten compared was chosen as a worst-performing estimate of the calibration density. For each subgroup size, a hundred trials of picking ten subgroups at that size was performed. Using these trials, a mean worst-performing estimate of the calibration density was calculated for each subgroup size, as well as standard error bars (figure 8).



#### 2.4.5. Tightness

It is desirable to score the quality of the uncertainty intervals irrespective of calibration. For example, assessing calibration might indicate that a small uncertainty interval could be sufficient to capture the in-domain ground truth across most predictions, but a UQ method might often be returning unnecessarily larger intervals in practice. The uncertainty intervals should only be as large as necessary to capture the ground truth, so scoring metrics assess what we herein refer to as the tightness of a UQ method. Tightness was measured by computing a negatively oriented mean interval score [40]. An increasingly lower score refers to an increasingly tight and appropriate uncertainty estimate, figure 2(c).

This mean interval score rewards itself for smaller uncertainty intervals that capture the ground truth but punishes itself for smaller uncertainty intervals that do not. Across all predictions, we drew uncertainty intervals across a ground-truth coverage of 1%–99% in 1% increments. The interval length corresponds to the coverage of our Gaussian distribution assumption (e.g.  $\mu \pm 1\sigma$  is a coverage of  $\sim 68\%$ ,  $\mu \pm 2\sigma$  is a coverage of  $\sim 95\%$ , etc). For each prediction, we computed a mean interval score across all coverage levels, figure 2(b). A mean interval score for the entire test set was computed by calculating a mean of these mean scores.

### 2.5. Scalar recalibration



The UQ methods tested were recalibrated [61] such that the uncertainty interval of  $\mu \pm 3\sigma$  more reliably behaves as a confidence interval and encloses the ground truth, which refers to the target property. An interval of  $\pm 3\sigma$  was chosen because such an interval provides 99.7% confidence under a perfect calibration assumption. Scalar recalibration was used, which involves multiplying all the uncertainty estimates  $\sigma$  across all adsorption energy predictions by a constant [57]. This recalibration method retains the original shape of the uncertainty distribution and stretches or shrinks the distribution based on scaling  $\sigma$ , so UQ methods with uncertainty distributions that reasonably fit our distribution-specific Gaussian assumption before recalibration are best recalibrated with this method. The constant was chosen via a black-box optimization algorithm that attempts to minimize the miscalibration area for each UQ method, table S4. Brent's method was used for this recalibration method because—unlike Platt scaling [62] for example—it is a nonparametric method that makes no assumptions about the shape of the calibration curve(s) to be recalibrated [63]. The Scikit-learn implementation of Brent's method was used [64].

## 3. Results and discussion

We assess the accuracy of each UQ method applied onto our CGCNN model by comparing parity plots and different quantitative metrics of accuracy, figure 3. The three UQ methods have similar distributions of adsorption energy predictions. By all quantitative measures, evidential regression outperforms the other UQ methods in accuracy. Evidential regression has an  $R^2$  of 0.902 (figure 3(c)), whereas MC dropout has 0.893 (figure 3(b)) and five-fold ensemble has 0.890 (figure 3(a)). Evidential regression has the lowest *M*DAE of 0.391 eV, suggesting that this UQ method is the most accurate on non-outlier data. Generally, five-fold ensemble and MC dropout perform similarly across all quantitative metrics of accuracy.

Although each UQ method (figures 3(a)–(c)) differs in its approach for estimating the uncertainty—ensembling involves segregating the training data, MC dropout involves modifying the model architecture, and evidential regression requires modifying both the model architecture and optimization—the accuracy is quite similar for all three methods. To explore why, we examine if the models are overfit or underfit. Consider MC dropout; this method is used to prevent model overfitting as specified by the dropout rate hyperparameter [65]. A higher dropout rate results in dropping out more nodes, thereby resulting in larger adjustments away from overfitting. However, we do not observe improvements in test accuracy as dropout rate increases, suggesting that the baseline CGCNN model that we apply UQ methods onto is not overfitting to the training data—but rather underfitting relative to the chemical space represented by the training data, figure S1. A plausible explanation for the underfitting of the three UQ methods with a CGCNN is that the OC20 dataset is very sparse; roughly 0.07% of possible calculations were performed when considering the dataset constraints on adsorbates, surfaces, and bulk compositions [34]. Consequently, our CGCNN model architecture may underfit the chemical space represented by the full IS2RE training dataset. To achieve higher accuracy, more state-of-the-art deep NN models and representations are needed, which is an area of on-going research [12].

We explored the sensitivity of the hyperparameter choice  $\lambda$  for evidential regression. Table S2 contains the test MAE results for evidential regression across varying values of hyperparameter  $\lambda$ . The test MAE across  $\lambda$  fluctuated at most by 0.48% from any given measurement, demonstrating that  $\lambda$  can be used to tune the size of the uncertainty intervals for this dataset and model without appreciable changes in the average accuracy of adsorption energy predictions. Based on this sensitivity analysis, we chose  $\lambda = 0.05$  for all




subsequent experiments because this regularization weight gives the most conservative uncertainty intervals of the non-zero weights tested, figure S3(b).

We construct violin box plots to assess both the dispersion and sharpness of each UQ method, as displayed in figure 4. High dispersion is desirable because one can define heuristics or rules based on the uncertainty to select and study catalysts for further study. For example, one might want to segregate very uncertain adsorbate/alloy systems from the rest for further study, but there may not be very uncertain systems to segregate if the dispersion is small. MC dropout is the most disperse UQ method on an absolute scale based on having the largest IQR. Five-fold smallest IQR; for this UQ method, defining uncertainty heuristics to select adsorbate/alloy systems for further study could be difficult—at least if one wants to select systems from the bulk of the uncertainty distribution. Evidential regression has dispersion in-between MC dropout and five-fold ensemble in terms of IQR, although we found the dispersion of evidential regression to vary substantially with varying values of  $\lambda$ , figure S3(b).

The  $C_v$  analysis demonstrates that the evidential regression results are the most disperse relative to the distribution mean. Evidential regression has a  $C_v$  of 2.13, whereas five-fold ensemble and MC dropout have a  $C_v$  of 0.786 and 0.779, respectively. Although MC dropout is the most disperse in absolute terms of the distribution spread, the dispersion of values around the distribution mean is poor. We conclude that evidential regression is the most disperse UQ method based on IQR and  $C_v$  collectively.

Of the methods considered, evidential regression is the least sharp (i.e. has the highest sharpness value); thus, for this dataset evidential regression gives the least confident adsorption energy predictions on average.




Uncertainty estimates of the adsorption energy predictions can range from 0 eV upwards for each UQ method, so we expected MC dropout to be the least sharp due to having the largest spread as measured by IQR, but this is not the case. We observed that evidential regression gives increasingly large—sometimes even enormous—uncertainty estimates for outliers as  $\lambda$  increases, which skews the sharpness. Whereas evidential regression is the least sharp (0.780 eV), five-fold ensemble is the sharpest (0.168 eV). Therefore, the five-fold ensemble gives the most confident adsorption energy predictions on average. We summarize IQR,  $C_v$ , and sharpness for each UQ method in table S1. Dispersion as measured by  $C_v$  for the UQ methods does not change after scalar recalibration. Before recalibration, MC dropout is the most disperse as measured by IQR, but this UQ method has poor dispersion as measured by  $C_v$ . While evidential regression does not have a better IQR before recalibration, this method is generally more disperse before and after recalibration when IQR and  $C_v$  are collectively considered. The effects of recalibration are discussed more below.

Although it is important for a UQ method to make confident predictions on average, one needs to ensure that the associated uncertainty estimate  $\sigma$  for each adsorption energy prediction are neither overconfident nor underconfident such that we can interpret the uncertainty estimate similarly to that of a confidence interval—an interval that reliably suggests where the actual ground-truth adsorption energy could exist. In other words, it is important to ensure that the UQ method is calibrated.

We assess average calibration across the test set in figure 5 before scalar recalibration. Each calibration curve has an associated miscalibration area, which refers to the area between any given calibration curve and the diagonal. A higher miscalibration area means worse model calibration (i.e. the reliability of using an uncertainty estimate as a confidence interval is weakened). The average calibration between the three UQ methods tested significantly varies, figure 5(a). The data shows that MC dropout provides the most calibrated adsorption energy predictions on average by having the smallest miscalibration area (0.20). Oppositely, five-fold ensemble gave the least calibrated adsorption energy predictions on average by having the largest miscalibration area (0.38). With a miscalibration area of 0.34, evidential regression performed similarly to the five-fold ensemble.

We observe a noticeable change in the degree of calibration for MC dropout with an increasing number of samples taken to construct the uncertainty estimate, figure 5(b). Five-sample uncertainty estimates are poorly calibrated, giving a miscalibration area of 0.36. However, 1000-sample uncertainty estimates reveal that MC dropout is the most calibrated UQ method tested before recalibration. The takeaway is straightforward; one needs to gather enough samples for the uncertainty estimates of MC dropout to converge. For our experimental setup, we found that a sample size of 50 nearly converges MC dropout, figure 5(b). Herein, we report the uncertainty of each adsorption energy prediction using a sample size of 1000 to ensure convergence.

The data in figure 5(c) demonstrates the effect of varying  $\lambda$  on the average calibration of evidential regression. Values of 0.0, 0.05, 0.1, 0.15, and 0.2 for  $\lambda$  were tested. Evidential regression becomes monotonically more calibrated with increasing  $\lambda$  for the values tested. For  $\lambda$  values above 0.2, the model struggled to converge onto a finite uncertainty estimate, which we discuss in SI. Interestingly, the calibration curves between all UQ methods were largely of a similar character—most remain below the diagonal. Because of the method of construction, calibration curves below the diagonal indicate overconfident adsorption energy predictions. It is known that increasing  $\lambda$  generally inflates the uncertainty estimate of a



**Figure 6.** Calibration curves to assess recalibration using scalar recalibration. Calibration curves on the test set are shown for (a) five-fold ensemble, (b) 1000-sample MC dropout, and (c) evidential regression ( $\lambda = 0.05$ ) before and after recalibration. The dashed red line denotes perfect average calibration. The miscalibration area before and after recalibration is shown inset.

**Table 1.** Mean interval score for each UQ method. A lower score indicates better tightness.

UQ method	Five-fold ensemble	1000-sample MC dropout	Evidential regression ( $\lambda = 0.05$ )
Interval score	5.307	4.131	3.901
Interval score (recalibrated)	2.907	3.472	3.550

prediction [51], so it is expected that the evidential regression calibration will improve as overconfident, small uncertainty estimates become appropriately inflated.


If a UQ method does not give calibrated and therefore reliable materials predictions as implemented, the method can be recalibrated. The results in figure 6 show the associated calibration curves on the test set before and after recalibrating each UQ method with scalar recalibration. In general, all three UQ methods saw a sizable decrease in average miscalibration area after recalibration. The data in figure 6(a) shows that five-fold ensembling gives the most calibrated adsorption energy predictions on average based on having the smallest miscalibration area. Despite the highest miscalibration area after recalibration, MC dropout as shown in figure 6(b) still gives competitively calibrated adsorption energy predictions on average. The calibration performance of evidential regression displayed in figure 6(c) is close to that of MC dropout. All three UQ methods give both overconfident and underconfident uncertainty estimates on average after recalibration, but the degree of miscalibration is less severe.

All UQ methods we trained have average calibration curves close to the diagonal for the test set after recalibration, which demonstrates that each UQ method is well-calibrated on average for many adsorption energy predictions. Having high calibration on average is useful because selecting adsorbate/alloy systems for further study based on the adsorption energy uncertainty (e.g. in active learning workflows) should be more reliable compared to if the model is highly miscalibrated.

Table 1 reports the mean interval score to assess the tightness of each UQ method before and after recalibration. Before recalibration, evidential regression gives the most appropriately tight uncertainty intervals on average despite not being the most calibrated UQ method. For the predictions that happen to be calibrated, the associated uncertainty intervals of evidential regression are the leanest on average. After recalibration, five-fold ensemble is the tightest UQ method on average.

UQ metrics holistically provide a means of comparing the shapes of uncertainty distributions across methods. For example, dispersion quantifies the spread of uncertainty, and sharpness quantifies a notion of average uncertainty. When considering these metrics in their entirety, no UQ method is clearly advantageous before recalibration. MC dropout is the most calibrated on average which might lead one to believe that this method is superior, but evidential regression outperforms the secondary criteria of accuracy, dispersion, and tightness. After recalibration, all UQ methods are comparably calibrated on average, but evidential regression outperforms in measures of accuracy and dispersion. We emphasize that the UQ metrics discussed provide a portfolio from which to holistically compare methods to robustly assess trustworthiness; any single metric on its own can give a misleading conclusion of trustworthiness.

Thus far, we have compared three UQ methods in terms of the five UQ metrics for our challenging benchmark, the OC20 dataset. To demonstrate how these modeling concepts can be used in practice, we report our uncertainty-guided enumeration of materials for a hydrogen adsorbate in figure 7. We subselect




609 hydrogen adsorbate/catalyst systems from the OC20 Val-ID test dataset used to assess each UQ method and filter these systems based on defined search criteria. Our search criteria for enumerating materials involves defining an adsorption energy range and an uncertainty limit to select systems.

For our case study, we select systems with hydrogen adsorption energies in the range of  $-0.1$ – $0.1$  eV. Such a range may be useful, for example, for studying materials for the hydrogen evolution reaction because hydrogen adsorption energy is a descriptor of catalytic activity [66, 67]. From these systems, we only choose those with an associated uncertainty  $\sigma$  of 0.05 eV or smaller. This adsorption energy range was chosen because each UQ method has inaccuracy due to the performance of the underlying CGCNN model. We decided to select a conservative, narrow search range to select only those predictions that should be confident. With our choice of  $\sigma$ , a  $\mu \pm 3\sigma$  interval implies that the ground-truth adsorption energy prediction should often be at most 0.15 eV away from the ML prediction—the reliability of which becomes less consistent with larger miscalibration. Any material satisfying these criteria is enumerated.

For this case study, our goal is not to rely on our ML adsorption energy predictions for the best chemical accuracy; rather, we demonstrate uncertainty estimates that are at least reliable and trustworthy enough to effectively enumerate materials. Additionally, from a modeling standpoint, we intend to demonstrate the increase in trustworthiness for predictions made with our model on the challenging OC20 dataset before and after recalibration.

We observe improved reliability in the uncertainty estimates after recalibration. Before recalibration, none of the UQ methods enumerated materials that matched our search criteria. After recalibration, evidential regression gave 15 catalyst materials matching the search criteria (figure 7). Materials that have the same bulk composition differ by adsorption site and Miller index. Crystallographic differences between materials such as  $\text{K}_2$  and  $\text{K}_8$  are given in table S5.

A UQ method should propose materials that match the search criteria for an application; however, these proposals should be honest. Six of these materials' predictions are honest in that their associated uncertainty interval of  $\mu \pm 3\sigma$  successfully captures the ground truth values. The rest of the proposed materials were dishonest. This lackluster overall screening performance is to be expected using this challengingly sparse OC20 dataset. First and foremost, the trained CGCNN model is not highly accurate as measured by MAE, so many of the adsorption energy predictions are far away from the ground truth values before even applying the UQ methods. A more accurate model would place the predictions closer to the ground truth, demonstrating the importance of recent efforts to develop even more accurate deep NN models [29, 33, 36, 37]. Given that the OC20 dataset is challenging, we note that the CGCNN architecture may have better accuracy on less sparse, narrowly selected datasets, as we observe in table S3. For evidential regression ( $\lambda = 0.05$ ), we report the element-wise sharpness and accuracy as measured by MAE for catalyst materials in figure S7, finding that periodic groups 3–5 are particularly uncertain and inaccurate. Additionally, we report



uncertainty distributions of evidential regression on a per-adsorbate basis in figure S8, finding that a handful of adsorbates (e.g.  $^*\text{NO}_2\text{NO}_2$ ) are particularly uncertain for this dataset, model, and UQ method. An improvement in model accuracy should plausibly make any given search criteria more forgiving—assuming that the UQ method applied onto the model is calibrated. We deduce this point from the effective calibration assumption that necessitates that the uncertainty  $\sigma$  of any prediction is probabilistically likely to grow or shrink with the accuracy of that prediction.

Nevertheless, it is noteworthy that we could identify trustworthy materials predictions as shown in figure 7, which highlights the importance of recalibration. From a purely data-driven viewpoint, we demonstrate that scalar recalibration did improve a UQ method's ability to enumerate DFT-accurate adsorption energies of these materials, although we acknowledge that many of these materials found in the OC20 dataset are not commonly used catalysts for reactions involving hydrogen. In a real catalyst screening study, one might consider additional criteria outside of adsorption energies, such as stability or selectivity.

Recalibration can be considered as a postprocessing step for a UQ method that can improve an initially poor result, which we demonstrate in this discussion. However, recalibration may not be necessary and may even introduce some drawbacks to the experimental design, depending on the method. For scalar recalibration, all uncertainty estimates  $\sigma$  are multiplied by a constant, which has the effect of multiplying the sharpness by said constant. Because sharpness is akin to the average uncertainty of a prediction, scalar recalibration scales the average magnitude of uncertainty, thereby making it more difficult to effectively define narrow  $\sigma$  screening criteria if the scaling enlarges the sharpness.

For our demonstration of enumerating hydrogen adsorbate/catalyst systems, we rationalize the increase in the number of enumerated systems after recalibration based on two reasons. First, each UQ method makes overconfident adsorption energy predictions, so inflating all the uncertainty estimates  $\sigma$  for any UQ method by a scalar made the model more appropriately calibrated—thereby improving the confidence interval reliability of each uncertainty estimate as supported by figure 7. Secondly, we define our uncertainty search criteria as  $\sigma = 0.05$  eV. Evidential regression ( $\lambda = 0.05$ ) shown in figure 4 is more forgiving for this screening criterion in that the uncertainty estimates  $\sigma$  are much less than 0.05 eV before recalibration such that inflating the uncertainty intervals did not exceed the  $\sigma = 0.05$  criterion after recalibration, unlike many systems for five-fold ensemble and MC dropout.

Although the average recalibration results of figure 6 appear excellent and the recalibration did improve the results of our material enumeration case study for the hydrogen adsorbate, a more nuanced analysis of recalibration on the scale of individual predictions was performed by estimating the calibration density. The adversarial group calibration results in figure 8 allow us to estimate the upper bound miscalibration associated with the calibration density across group sizes. We only report here adversarial calibration of group sizes up to 2% of the test set size because we otherwise observe asymptotic miscalibration.

Generally, the miscalibration across UQ methods monotonically improves across group sizes after recalibrating. This consistency in behavior is true near the limiting case of a group size of zero, that is,

individual calibration of predictions. However, directly assessing individual calibration is often unverifiable for finite dataset sizes, despite being a strict calibration constraint in the literature [61].

Although we previously calculate average calibration curves across the entire test set, a UQ method is truly calibrated if all individual predictions are themselves calibrated. The recalibration results demonstrate significant improvement in the average calibration (i.e. calibration corresponding to a 100% group size), yet the limiting behavior near the origin in figure 8 suggests that there is only a mild improvement in the individual calibration for the most miscalibrated predictions. Ultimately, the analysis in figure 8 helps rationalize why many of the uncertainty intervals in the figure 7 demonstration do not capture the ground truth. Moreover, we highlight a subtlety that the dramatic average recalibration improvements shown in figure 6 are misleading and, in actuality, more modest, so the interpretation of average calibration needs to be handled carefully for applications.

## 4. Conclusions

We compare three UQ methods to frame future analyses in high-throughput materials discovery by training a state-of-the-art CGCNN on the challenging OC20 dataset. These methods were chosen for comparison based on their general applicability to a wide range of NN architectures across materials problems. We base our comparison on the UQ metrics, which quantify not only the accuracy of predictions but also the size, spread, and reliability of the associated uncertainty intervals. These intervals—if they are trustworthy—can be used to infer the accuracy of unseen materials predictions in high-throughput discovery schemes. Before recalibration, evidential regression is advantageously found to be the most accurate, disperse, and tight UQ method, but MC dropout is the most calibrated method on average. After recalibration, evidential regression is found to be the most accurate and disperse, as well as competitively calibrated. Additionally, evidential regression provides high utility by enabling tunable uncertainty estimates that output upon prediction time, unlike  $k$ -fold ensembles and MC dropout. This tunability allows for more flexibility in addressing application-specific shortcomings with sharpness, dispersion, and tightness. For future high-throughput studies using distribution-specific UQ, we recommend evidential regression because of its tunability, demonstrated trustworthiness, and computational tractability on this challenging dataset benchmark.

Through the enumeration of OC20 materials in the case of a hydrogen adsorbate, scalar recalibration demonstrably enhances the trustworthiness of evidential regression. We note that scalar recalibration was most effective on the tested UQ methods that give uncertainty distributions of a similar shape to our Gaussian-specific assumption. We conjecture based on the experimental results that effective recalibration is a promising knowledge gap to make distribution-specific UQ more accessible across materials studies. NNs can perform surprisingly poorly under domain shift [68], so it is unclear how each UQ method would perform across different model architectures and complex material spaces. With effective recalibration that has little negative effect on the UQ metrics, the choice of UQ method would be more arbitrary. Researchers would be able to prioritize choosing the method that is the most computationally tractable and recalibrate accordingly. The UQ metrics and analysis discussed can serve as a frame of reference for this endeavor.

## Data availability statement

The data that support the findings of this study are openly available at the following URL/DOI: <https://github.com/CGruich/ClarifyingTrustUQ>. Data will be available from 2 June 2023.

## Acknowledgments

C G acknowledges support from the NSF Graduate Research Fellowship Program. B R G acknowledges support from the MICDE Catalyst Grant from the Michigan Institute for Computational Discovery and Engineering. This research used the National Energy Research Scientific Computing Center resources, a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE 1841052.

## Associated content

The supporting information is available free of charge on the journal. Data and code supporting the results of this study are available through the following URL: <https://github.com/CGruich/ClarifyingTrustUQ>. The supporting information contains details on UQ model dispersion and sharpness; the effect of dropout rate

on the CGCNN; aleatoric and epistemic uncertainty analysis for hyperparameter  $\lambda$ ; supplementary adversarial group calibration analysis; material/adsorbate outliers.

### Author contributions

Conceptualization, C G, B R G; methodology, C G; software, C G; formal analysis, C G, V M; investigation, C G, V M; resources, C G, B R G; data curation, C G; writing—original draft, C G; writing—review & editing, C G, V M, Y W, B R G; visualization, C G; supervision, B R G; project administration, B R G; funding acquisition, B R G.

### Conflict of interest

The authors declare no competing financial interest.

### ORCID iDs

Cameron J Gruich  <https://orcid.org/0000-0002-3801-1296>

Varun Madhavan  <https://orcid.org/0000-0002-5107-4793>

Yixin Wang  <https://orcid.org/0000-0002-6617-4842>

### References

- [1] Goldsmith B R, Esterhuizen J, Liu J, Bartel C J and Sutton C 2018 Machine learning for heterogeneous catalyst design and discovery *AIChE J.* **64** 2311–23
- [2] Schleder G R, Padilha A C M, Acosta C M, Costa M and Fazzio A 2019 From DFT to machine learning: recent approaches to materials science—a review *J. Phys. Mater.* **2** 032001
- [3] Schlexer Lamoureux P, Winther K T, Garrido Torres J A, Streibel V, Zhao M, Bajdich M, Abild-Pedersen F and Bligaard T 2019 Machine learning for computational heterogeneous catalysis *ChemCatChem* **11** 3581–601
- [4] Yang W, Fidelis T T and Sun W-H 2020 Machine learning in catalysis, from proposal to practicing *ACS Omega* **5** 83–88
- [5] Zhong M *et al* 2020 Accelerated discovery of CO<sub>2</sub> electrocatalysts using active machine learning *Nature* **581** 178–83
- [6] Tran K and Ulissi Z W 2018 Active learning across intermetallics to guide discovery of electrocatalysts for CO<sub>2</sub> reduction and H<sub>2</sub> evolution *Nat. Catal.* **1** 696–703
- [7] Tran R, Wang D, Kingsbury R, Palizhati A, Persson K A, Jain A and Ulissi Z W 2022 Screening of bimetallic electrocatalysts for water purification with machine learning *J. Chem. Phys.* **157** 074102
- [8] Abdar M *et al* 2021 A review of uncertainty quantification in deep learning: techniques, applications and challenges *Inf. Fusion* **76** 243–97
- [9] Meredig B 2019 Five high-impact research areas in machine learning for materials science *Chem. Mater.* **31** 9579–81
- [10] Hie B, Bryson B D and Berger B 2020 Leveraging uncertainty in machine learning accelerates biological discovery and design *Cell Syst.* **11** 461–77.e9
- [11] Walker E A 2022 Bayesian statistics in catalysis: a perspective *Curr. Opin. Chem. Eng.* **36** 100820
- [12] Kolluru A, Shuaibi M, Palizhati A, Shoghi N, Das A, Wood B, Zitnick C L, Kitchin J R and Ulissi Z W 2022 Open challenges in developing generalizable large scale machine learning models for catalyst discovery *ACS Catal.* **12** 8572–81
- [13] Gal Y and Ghahramani Z 2016 Dropout as a Bayesian approximation: representing model uncertainty in deep learning *Proc. of the 33rd Int. conf. on machine learning* vol 48 pp 1050–9
- [14] Tran K, Neiswanger W, Yoon J, Zhang Q, Xing E and Ulissi Z W 2020 Methods for comparing uncertainty quantifications for material property predictions *Mach. Learn.: Sci. Technol.* **1** 025006
- [15] Hu Y, Musielewicz J, Ulissi Z W and Medford A J 2022 Robust and scalable uncertainty estimation with conformal prediction for machine-learned interatomic potentials *Mach. Learn.: Sci. Technol.* **3** 045028
- [16] Janet J P, Duan C, Yang T, Nandy A and Kulik H J 2019 A quantitative uncertainty metric controls error in neural network-driven chemical discovery *Chem. Sci.* **10** 7913–22
- [17] Noh J, Gu G H, Kim S and Jung Y 2020 Uncertainty-quantified hybrid machine learning/density functional theory high throughput screening method for crystals *J. Chem. Inf. Model* **60** 1996–2003
- [18] Wander B, Broderick K and Ulissi Z W 2022 Atlas: an automated framework for catalyst discovery demonstrated for direct syngas conversion *Catal. Sci. Technol.* **12** 6256–67
- [19] Hutchinson M L, Antono E, Gibbons B M, Paradiso S, Ling J and Meredig B 2017 Overcoming data scarcity with transfer learning (arXiv:1711.05099v1)
- [20] Settles B 2012 *Active Learning* (Cham: Springer) (<https://doi.org/10.1007/978-3-031-01560-1>)
- [21] Kim M, Ha M Y, Jung W, Yoon J, Shin E, Kim I, Lee W B, Kim Y and Jung H 2022 Searching for an optimal multi-metallic alloy catalyst by active learning combined with experiments *Adv. Mater.* **34** 2108900
- [22] Tang Q, Lau Y B, Hu S, Yan W, Yang Y and Chen T 2010 Response surface methodology using Gaussian processes: towards optimizing the trans-stilbene epoxidation over Co<sup>2+</sup>-NaX catalysts *Chem. Eng. J.* **156** 423–31
- [23] Ulissi Z W, Singh A R, Tsai C and Nørskov J K 2016 Automated discovery and construction of surface phase diagrams using machine learning *J. Phys. Chem. Lett.* **7** 3931–5
- [24] Mamun O, Winther K T, Boes J R and Bligaard T 2020 A Bayesian framework for adsorption energy prediction on bimetallic alloy catalysts *npj Comput. Mater.* **6** 177
- [25] Pedersen J K *et al* 2021 Bayesian optimization of high-entropy alloy compositions for electrocatalytic oxygen reduction *Angew. Chem., Int. Ed.* **60** 24144–52
- [26] Williams C K I and Rasmussen C E 1995 Gaussian processes for regression *Advances in Neural Information Processing Systems* pp 514–20



- [27] Zhan N and Kitchin J R 2022 Model-specific to model-general uncertainty for physical properties *Ind. Eng. Chem. Res.* **61** 8368–77
- [28] Liu H, Ong Y-S, Shen X and Cai J 2020 When Gaussian process meets big data: a review of scalable GPs *IEEE Trans. Neural Netw. Learn. Syst.* **31** 4405–23
- [29] Qiao Z, Welborn M, Anandkumar A, Manby F R and Miller I I T F 2020 OrbNet: deep learning for quantum chemistry using symmetry-adapted atomic-orbital features *J. Chem. Phys.* **153** 124111
- [30] Zitnick C L, Das A, Kolluru A, Lan J, Shuaibi M, Sriram A, Ulissi Z and Wood B 2022 Spherical channels for modeling atomic interactions (arXiv:2206.14331v2)
- [31] Gasteiger J, Shuaibi M, Sriram A, Günnemann S, Ulissi Z, Zitnick C L and Das A 2022 How do graph networks generalize to large and diverse molecular systems? (arXiv:2204.02782v1)
- [32] Liao Y-L and Smidt T 2022 Equiformer: equivariant graph attention transformer for 3D atomistic graphs (arXiv:2206.11990v2)
- [33] Brandstetter J, Hesselink R, van der Pol E, Bekkers E J and Welling M 2022 Geometric and physical quantities improve E(3) equivariant message passing (arXiv:2110.02905v3)
- [34] Chanussot L *et al* 2021 The Open Catalyst 2020 (OC20) dataset and community challenges *ACS Catal.* **11** 6059–72
- [35] Xie T and Grossman J C 2018 Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties *Phys. Rev. Lett.* **120** 145301
- [36] Schütt K T, Sauceda H E, Kindermans P-J, Tkatchenko A and Müller K-R 2018 SchNet—a deep learning architecture for molecules and materials *J. Chem. Phys.* **148** 241722
- [37] Gasteiger J, Giri S, Margraf J T and Günnemann S 2022 Fast and uncertainty-aware directional message passing for non-equilibrium molecules (arXiv:2011.14115v3)
- [38] Schaarschmidt M, Riviere M, Ganose A M, Spencer J S, Gaunt A L, Kirkpatrick J, Axelrod S, Battaglia P W and Godwin J 2022 Learned force fields are ready for ground state catalyst discovery (arXiv:2209.12466v1)
- [39] Gneiting T, Balabdaoui F and Raftery A E 2007 Probabilistic forecasts, calibration and sharpness *J. R. Stat. Soc. B* **69** 243–68
- [40] Gneiting T and Raftery A E 2007 Strictly proper scoring rules, prediction, and estimation *J. Am. Stat. Assoc.* **102** 359–78
- [41] Gneiting T and Katzfuss M 2014 Probabilistic forecasting *Annu. Rev. Stat. Appl.* **1** 125–51
- [42] Guo C, Pleiss G, Sun Y and Weinberger K Q 2017 On calibration of modern neural networks *Int. Conf. on Machine Learning* vol 70 pp 1321–30
- [43] Levi D, Gispan L, Giladi N and Fetaya E 2022 Evaluating and calibrating uncertainty prediction in regression tasks *Sensors* **22** 5540
- [44] Pernet P 2022 The long road to calibrated prediction uncertainty in computational chemistry *J. Chem. Phys.* **156** 114109
- [45] Michelmore R, Wicker M, Laurenti L, Cardelli L, Gal Y and Kwiatkowska M 2020 Uncertainty quantification with statistical guarantees in end-to-end autonomous driving control 2020 *IEEE Int. Conf. on Robotics and Automation (ICRA)* (Paris: IEEE) pp 7344–50
- [46] Kwon Y, Won J-H, Kim B J and Paik M C 2020 Uncertainty quantification using Bayesian neural networks in classification: application to biomedical image segmentation *Comput. Stat. Data Anal.* **142** 106816
- [47] Olivier A, Shields M D and Graham-Brady L 2021 Bayesian neural networks for uncertainty quantification in data-driven materials modeling *Comput. Methods Appl. Mech. Eng.* **386** 114079
- [48] Dietterich T G 2000 Ensemble methods in machine learning *Multiple Classifier Systems Lecture Notes in Computer Science* vol 1857 (Berlin: Springer) pp 1–15
- [49] Lakshminarayanan B, Pritzel A and Blundell C 2017 Simple and scalable predictive uncertainty estimation using deep ensembles *Advances in Neural Information Processing Systems* **31** (<https://doi.org/10.5555/3295222.3295387>)
- [50] Liu J Z, Kioumourtzoglou M-A, Paisley J and Coull B A 2019 Accurate uncertainty estimation and decomposition in ensemble learning 33rd *Int. conf. Advances in Neural Information Processing Systems* vol 12 pp 8952–63
- [51] Amini A, Schwarting W, Soleimany A and Rus D 2020 Deep evidential regression 34th *Int. conf. Advances in Neural Information Processing Systems* vol 19 pp 14927–37
- [52] Soleimany A P, Amini A, Goldman S, Rus D, Bhatia S N and Coley C W 2021 Evidential deep learning for guided molecular property prediction and discovery *ACS Cent. Sci.* **7** 1356–67
- [53] Humphreys J, Lan R and Tao S 2021 Development and recent progress on ammonia synthesis catalysts for Haber–Bosch process *Adv. Energy Sustain. Res.* **2** 2000043
- [54] Sadrameli S M 2016 Thermal/catalytic cracking of liquid hydrocarbons for the production of olefins: a state-of-the-art review II: catalytic cracking review *Fuel* **173** 285–97
- [55] Rahimpour M R, Jafari M and Iranshahi D 2013 Progress in catalytic naphtha reforming process: a review *Appl. Energy* **109** 79–93
- [56] Chung Y and Neiswanger W 2021 Beyond pinball loss: quantile methods for calibrated uncertainty quantification *Advances in Neural Information Processing Systems* vol 34 pp 10971–84 (arXiv:2011.09588v4)
- [57] Chung Y, Char I, Guo H, Schneider J and Neiswanger W 2021 Uncertainty toolbox: an open-source library for assessing, visualizing, and improving uncertainty quantification (arXiv:2109.10254v1)
- [58] Hüllermeier E and Waegeman W 2021 Aleatoric and epistemic uncertainty in machine learning: an introduction to concepts and methods *Mach. Learn.* **110** 457–506
- [59] Scott D W 1992 *Multivariate Density Estimation: Theory, Practice, and Visualization* (New York: Wiley) (<https://doi.org/10.1002/9780470316849>)
- [60] Kuleshov V, Fenner N and Ermon S 2018 Accurate uncertainties for deep learning using calibrated regression 35th *Int. Conf. on Machine Learning* vol 80 pp 2796–804 (arXiv:1807.00263v1)
- [61] Zhao S, Ma T and Ermon S 2020 Individual calibration with randomized forecasting 37th *Int. Conf. on Machine Learning* vol 119 pp 11387–97
- [62] Platt J C 2000 Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods *Adv. Large Margin Classif.* **10** 61–74
- [63] Brent R P 1971 An algorithm with guaranteed convergence for finding a zero of a function *Comput J.* **14** 422–5
- [64] Pedregosa F *et al* 2011 Scikit-learn: machine learning in python *J. Mach. Learn. Res.* **12** 2825–30
- [65] Srivastava N, Hinton G, Krizhevsky A, Sutskever I and Salakhutdinov R 2014 Dropout: a simple way to prevent neural networks from overfitting *J. Mach. Learn. Res.* **30** 1929–58
- [66] Nørskov J K, Bligaard T, Logadottir A, Kitchin J R, Chen J G, Pandelov S and Stimming U 2005 Trends in the exchange current for hydrogen evolution *J. Electrochem. Soc.* **152** J23
- [67] Mohammed-Ibrahim J and Sun X 2019 Recent progress on earth abundant electrocatalysts for hydrogen evolution reaction (HER) in alkaline medium to achieve efficient water splitting—a review *J. Energy Chem.* **34** 111–60
- [68] Quiñero-Candela J, Sugiyama M, Schwaighofer A and Lawrence N 2009 *Dataset Shift in Machine Learning* (Cambridge, MA: MIT Press) (<https://doi.org/10.7551/mitpress/9780262170055.001.0001>)