

UC Berkeley

UC Berkeley Previously Published Works

Title

A critical examination of compound stability predictions from machine-learned formation energies

Permalink

<https://escholarship.org/uc/item/52r1r2bm>

Journal

npj Computational Materials, 6(1)

ISSN

2057-3960

Authors

Bartel, Christopher J
Trewartha, Amalie
Wang, Qi
[et al.](#)

Publication Date

2020

DOI

10.1038/s41524-020-00362-y

Peer reviewed

A critical examination of compound stability predictions from machine-learned formation energies

Christopher J. Bartel^{1*}, Amalie Trewartha¹, Qi Wang², Alex Dunn^{1,2}, Anubhav Jain², Gerbrand Ceder^{1,3*}

¹Department of Materials Science & Engineering, University of California, Berkeley, Berkeley, CA 94720, USA

²Energy Technologies Area, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA

³Materials Sciences Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA

*correspondence to cbartel@berkeley.edu, gceder@berkeley.edu

Abstract

Machine learning has emerged as a novel tool for the efficient prediction of materials properties, and claims have been made that machine-learned models for the formation energy of compounds can approach the accuracy of Density Functional Theory (DFT). The models tested in this work include five recently published compositional models, a baseline model using stoichiometry alone, and a structural model. By testing seven machine learning models for formation energy on stability predictions using the Materials Project database of DFT calculations for 85,014 unique chemical compositions, we show that while formation energies can indeed be predicted well, all compositional models perform poorly on predicting the stability of compounds, making them considerably less useful than DFT for the discovery and design of new solids. Most critically, in sparse chemical spaces where few stoichiometries have stable compounds, only the structural model is capable of efficiently detecting which materials are stable. The non-incremental improvement of structural models compared with compositional models is noteworthy and encourages the use of structural models for materials discovery. This work demonstrates that accurate predictions of formation energy do not imply accurate predictions of stability, emphasizing the importance of assessing model performance on stability predictions, for which we provide a set of publicly available tests.

Introduction

Machine learning (ML) is emerging as a novel tool for rapid prediction of material properties.¹⁻⁶ In general, these predictions are made by fitting statistical models on a large number of data points. Because of the scarcity of well-curated experimental data in materials science, this input data is often obtained from Density Functional Theory (DFT) calculations housed in one of the many open materials databases.⁷⁻¹² In principle, once these models are trained on this immense set of quantum chemical data, the determination of properties for new materials can be made in orders-of-magnitude less time using the trained models compared to computationally expensive DFT calculations.

Of particular interest is the use of machine learning to discover new materials. The combinatorics of materials discovery make for an immensely challenging problem – if we consider the possible combinations of just four elements (A, B, C, D), from any of the ~ 80 elements that are technologically relevant, there are already ~ 1.6 million quaternary chemical spaces to consider. This is before we consider such factors as stoichiometry ($ABCD_2, AB_2C_3D_4$, etc.) or crystal structure, each of which add substantially to the combinatorial complexity. The Inorganic Crystal Structure Database (ICSD) of known solid-state materials contains $\sim 10^5$ entries,¹³ several orders of magnitude less than the 10^{10} quaternary compositions identified as plausible using electronegativity- and charge-based rules.¹⁴ This suggests that 1) there is ample opportunity for new materials discovery and 2) the problem of finding stable materials may resemble the needle-in-a-haystack problem, with many unstable compositions for each stable one. The immensity of this problem is a natural fit for high-throughput machine learning techniques.

In this work, we closely examine whether recently published machine learning models for formation energy are capable of distinguishing the relative stability of chemically similar materials and provide a roadmap for doing the same for future models. We show that although the formation energy of compounds from elements can be learned with high accuracy using a variety of machine learning approaches, these learned formation energies do not reproduce DFT-calculated relative stabilities. While the accuracy of these models for formation energy approaches the DFT error (relative to experiment), DFT predictions benefit from a systematic cancellation of error^{15,16} when making stability predictions, while ML models do not. Of particular concern for most ML models is the high rate of materials predicted to be stable that are not stable by DFT, impeding the use of

these models to efficiently discover new materials. As a result, we suggest that more critical evaluation methods for machine learning methods are developed.

Results and Discussion

The relationship between formation energy and stability

A necessary condition for a material to be used for any application is stability (under some conditions). The thermodynamic stability of a material is defined by its Gibbs energy of decomposition, ΔG_d , which is the Gibbs formation energy, ΔG_f , of the specified material relative to all other compounds in the relevant chemical space. Temperature-dependent thermodynamics are not yet tractable with high-throughput DFT and have only sparsely been addressed with ML,¹⁷ so material stability is primarily assessed using the decomposition enthalpy, ΔH_d , which is approximated as the total energy difference between a given compound and competing compounds in a given chemical space.^{15,16,18,19} For the purpose of this study we will directly compare ML predictions and DFT calculations of ΔH_d , hence the lack of entropy contributions is not an issue.

The quantity ΔH_d is obtained by a convex hull construction in formation enthalpy (ΔH_f)-composition space. **Figure 1a** shows an example for a binary A - B space, having three known compounds, A_4B , A_2B , and AB_3 . The convex hull is the lower convex enthalpy envelope which lies below all points in the composition space (blue line). Stable compositions lie on the convex hull, and unstable compositions lie above the hull. A_4B is *unstable* (above the hull), so $\Delta H_d > 0$ and is calculated as the distance in ΔH_f between A_4B and the convex hull of stable points. AB_3 is stable (on the hull), so $\Delta H_d < 0$ and is calculated as the distance in ΔH_f between AB_3 and a hypothetical convex hull constructed without AB_3 (dashed line). $|\Delta H_d|$ is therefore the minimum amount that ΔH_f must decrease for an unstable compound to become thermodynamically stable or the maximum amount that ΔH_f can increase for a stable compound to remain stable. We used ΔH_d in this work instead of the more common, “energy above the hull”, because the former provides a distribution of values for stable compounds whereas the latter is equal to 0 for all stable compounds. The convex hull construction, described here for a binary system, generalizes for chemical spaces comprised of any number of elements.

Hence, while ΔH_f quantifies to what extent a compound may form from its elements, the thermodynamic quantity that controls phase stability is ΔH_d and arises from the competition between ΔH_f for all compounds within a chemical space. As we show later, while formation

enthalpies can be of the order of several eV the value of ΔH_d is typically 1-2 orders of magnitude smaller. In addition, thermodynamic stability is highly nonlinear in ΔH_d around zero, as negative values indicate stable compounds, whereas positive values are unstable or metastable compounds. Although ΔH_d determines stability, the standard thermodynamic property that is predicted by ML models is the absolute ΔH_f .^{20–29} This is in large part because ΔH_f is intrinsic to a given compound, whereas ΔH_d inherently depends upon a compound’s stability relative to neighboring compositions, making ΔH_d less robust to learn directly.

Using data available in the Materials Project (MP),¹⁶ we applied the convex hull construction to obtain ΔH_d for 85,014 inorganic crystalline solids (the majority of which are in the ICSD) and compare ΔH_d to ΔH_f in **Figure 1b**. It is clear that effectively no linear correlation exists between ΔH_d and ΔH_f , except for the trivial case where only a single compound exists in a chemical space ($\Delta H_d = \Delta H_f$), which is true for only $\sim 3\%$ of materials in MP. While ΔH_f somewhat uniformly spans a wide range of energies (mean \pm average absolute deviation = -1.42 ± 0.95 eV/atom), ΔH_d spans much smaller energies (0.06 ± 0.12 eV/atom), suggesting ΔH_d is the more sensitive or subtle quantity to predict (a histogram of ΔH_f and ΔH_d is provided in **Figure S1**). Still, while no linear correlation exists between ΔH_d and ΔH_f , and ΔH_d occurs over a much smaller energy range, it would be possible for ΔH_f models to predict ΔH_d as long as the relative differences in ΔH_f within a given chemical space are predicted with accuracy comparable to the range of variation in ΔH_d or if they would benefit from substantial error cancellation when comparing the energies of compounds with similar chemistry.

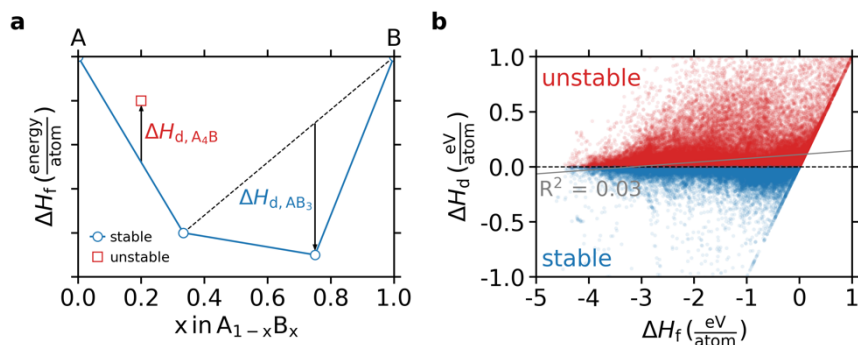


Figure 1. **a)** Illustration of the convex hull construction to obtain the decomposition enthalpy, ΔH_d , from the formation enthalpy, ΔH_f . **b)** The decomposition enthalpy, ΔH_d , shown against the formation enthalpy, ΔH_f , for 85,014 ground-state entries in Materials Project, indicating effectively no correlation between the two quantities. The strong linear correlation that is present at $\Delta H_d = \Delta H_f$ arises for chemical spaces that contain only one compound ($\sim 3\%$ of compounds). These compounds were excluded from the correlation coefficient, R^2 , determination.

Learning formation energy from chemical composition

Machine learning material properties requires that an arbitrary material is “represented” by a set of attributes (features). This representation can be as simple as a vector corresponding to the fractional amount of each element in the compound (e.g., $\text{Li}_2\text{O} = [0, 0, 2/3, 0, 0, 0, 0, 1/3, 0, 0, \dots]$, where the length of the vector is the number of elements in the periodic table), or a vector that includes substantial physical or chemical information about the material. In the search for new materials, the structure is rarely known *a priori*, and instead a list of compositions with unknown structure is screened for stability, i.e., the possibility that a thermodynamically stable structure exists at that composition. In this case, the material representation is constructed only from the chemical formula without including properties such as the geometric (i.e., lattice and basis) or electronic structure. These models, which take as input the chemical formula and output thermodynamic predictions, are henceforth referred to as *compositional models* here.

In this work, we assessed the potential for five recently introduced compositional models – *Meredig*,²⁰ *Magpie*,²¹ *AutoMat*,²² *ElemNet*,²³ and *Roost*²⁴ – to predict the stability of compounds in MP. *Meredig*, *Magpie*, and *AutoMat* include chemical information for each element in their material representations from quantities such as atomic electronegativities, radii, and elemental group. Each of these compositional models were trained using gradient-boosted regression trees (XGBoost³⁰). *ElemNet* and *Roost* differ in that no *a priori* information other than the stoichiometry is used as input. For *ElemNet*, a deep learning architecture maps the stoichiometry input into formation energy predictions. For *Roost*, the stoichiometric representation and fit are simultaneously learned using a graph neural network. In addition, we included a baseline model for comparison, *ElFrac*, where the representation is simply the stoichiometric fraction of each element in the formula, trained using XGBoost³⁰. Because compositional models necessarily make the same prediction for all structures having the same formula, all analysis in this work was performed using the lowest energy (ground-state) structure in MP for each available composition. Additional details on the training of each model and the MP dataset is available in the **Methods** section.

Parity plots comparing ΔH_f in MP ($\Delta H_{f,\text{MP}}$) to machine-learned ΔH_f ($\Delta H_{f,\text{pred}}$) for each model are shown in **Figure 2**. It is clear that each published representation substantially improves upon the baseline *ElFrac* model, decreasing the mean absolute error (MAE) by 27-74%. This increased accuracy is attributed to the increased complexity of the representation. For most models,

the MAE between MP and these ML models is comparable to the expected numerical disagreement between MP and experimentally obtained ΔH_f ,^{8,16,31–33} implying a substantial amount of the information required to determine ΔH_f is contained in the composition (and not the structure). The success of ML models for predicting ΔH_f is not surprising considering the historical context of simple heuristics that perform relatively well at predicting the driving force for the formation of compounds from elements – e.g., the Miedema model.³⁴




Figure 2. Parity plot for formation enthalpy predictions using six different machine learning models that take as input the chemical formula and output ΔH_f . *ElFrac* refers to a baseline model that parametrizes each formula only by the stoichiometric coefficient of each element. *Meredig*, *Magpie*, *AutoMat*, *ElemNet*, and *Roost* refer to Refs. ^{20–24}, respectively. $\Delta H_{f,\text{pred}}$ corresponds with ML predictions for aggregated hold-out sets during five-fold cross-validation of the Materials Project dataset (see **Methods** for details). $\Delta H_{f,\text{MP}}$ refers to the formation energy per atom in the MP database. The absolute error on ΔH_f is shown as the colorbar and the mean absolute error (MAE) is shown within each panel.

Implicit stability predictions from learned formation enthalpies

While the mean absolute error of the ML-predicted ΔH_f approaches the mean absolute error between DFT and experiment for this quantity ($\sim 0.1\text{--}0.2$ eV/atom)^{8,16,31,32,35}, the use of ΔH_f for stability predictions requires that the *relative* ΔH_f between chemically similar compounds is

predicted more accurately. To assess the accuracy of the relative ΔH_f , we reconstructed, for each ML model, the convex hulls for all chemical spaces using $\Delta H_{f,\text{pred}}$. Parity plots for ΔH_d are shown in **Figure 3**. Even though the quantity ΔH_d is on average much smaller than ΔH_f , the MAE in predicting it is almost identical to the error in predicting ΔH_f (**Figure 2**), indicating very little error cancellation for the ML models when energy differences are taken in a chemical space, which is in contrast to the beneficial error cancellation for stability predictions with DFT.^{15,16} In contrast to ΔH_f , where all representations substantially improve the predictive accuracy from the baseline *ElFrac* model, for ΔH_d , four of the five models (all except *Roost*) only negligibly improve upon the baseline model with MAE of ~ 0.11 - 0.14 eV/atom. Importantly, for the purposes of predicting stability, a difference of ~ 0.1 eV/atom can be the difference between a compound that is readily synthesizable and one that is unlikely to ever be realized.^{36,37}

DFT calculations benefit from a systematic cancellation of errors that leads to much smaller errors for ΔH_d than for ΔH_f , with MAE for ΔH_d as low as ~ 0.04 eV/atom for a substantial fraction of decomposition reactions.¹⁶ Unfortunately, ML models do not similarly benefit from this cancellation of errors and instead appear to learn clusters in material space that have similar ΔH_f , but they are generally unable to distinguish between stable and unstable compounds *within* a chemical space. It is notable that *Roost* substantially improves upon the other models. However, there are still strong signatures of inaccurate stability predictions in its parity plot (**Figure 3**), most notably in the \sim vertical line at $\Delta H_{d,\text{pred}} = 0$ and \sim horizontal line at $\Delta H_{d,\text{MP}} = 0$. These two lines indicate substantial disagreement between the actual and predicted stabilities for many compounds, despite the relatively low MAE.




Figure 3. Parity plot for decomposition enthalpy predictions. $\Delta H_{d,MP}$ results from convex hulls constructed with $\Delta H_{f,MP}$ (Fig. 2). $\Delta H_{d,pred}$ is obtained from convex hulls constructed with $\Delta H_{f,pred}$ (Fig. 2). The annotations are the same as in Fig. 2.

The inability for compositional models to properly distinguish relative stability is further demonstrated by assessing how well the models classify compounds as stable (on the convex hull) or unstable (above the hull), as shown in Figure 4. 60% of the compounds in the MP dataset are not on the hull, so the classification accuracy of a naïve model that states that everything is unstable would be 60%. Five of the six models (all except *Roost*) only marginally improve upon this extremely naïve model accuracy (58-65%). Strikingly, *Roost* considerably outperforms the other compositional models (76% accuracy), despite using stoichiometry alone as input. Plausibly, this superior performance is due to the use of weighted soft attention mechanisms during training of the representation.³⁸ Although only the nominal chemical composition (element fraction) is used as input, the model learns a more meaningful representation of this input composition on a case-by-case basis during training. This is in contrast to the other compositional models, which have fixed stoichiometric representations and either include hand-picked elemental attributes such as electronegativity (*Meredig* and *Magpie*) or use deep learning (*ElemNet*). Notably, *AutoMat* uses a two-step process: first it rationally selects the most relevant elemental attributes from a large list

using a decision tree model and then fits a regression model with the reduced feature space. Considering the modest classification accuracy by *AutoMat* (65%), despite the wide range of elemental attributes considered in its optimization, we speculate that further improvements in the clever selection of these attributes is unlikely to lead to transformative improvements in predicted stabilities. Instead, major improvements to compositional formation energy models will likely result from qualitative changes in model architecture, as in *Roost*, and not from optimizing the selection of elemental attributes.

While *Roost* improves considerably upon other compositional models, the accuracy, F_1 score, and false positive rate taken together do not inspire much confidence that any of these models can accurately predict the stability of solid-state materials (**Figure 4**). Of particular concern is the high false positive rates of 25-38%. This metric provides the likelihood that a compound predicted to be stable will not actually be. Further aggravating this situation is that the false positive rate reported here for the models is greatly underestimated compared to the false positive rate that is expected for new materials discovery. The MP database is largely populated with known materials extracted from the ICSD, and this results in $\sim 40\%$ of the entries in MP being on the hull. The fraction of all conceivable hypothetical materials (from which new materials will be discovered) that are stable is likely several orders of magnitude lower than 40%. This necessitates that searches for new materials cover a huge number of possible compounds, and false positive rates in excess of 25% would lead to an enormous amount of predicted materials which are not stable, limiting the ability for these ML models to efficiently accelerate new materials discovery.

A key consideration when discussing the accuracy in classifying compounds as stable or unstable is the choice of threshold for stability, which we have chosen to be $\Delta H_d = 0$. In materials discovery or screening applications, compounds are often considered potentially synthesizable even for $\Delta H_d > 0$ to consider potential inaccuracies in the predicted stabilities and account for a range of accessible metastability.^{36,37} To probe the effects of moving this threshold for stability to higher or lower values of ΔH_d , we show the receiver operating characteristic (ROC) curves for each model in **Figure S2**. As the threshold for stability moves to larger positive values, increases in model accuracy are concomitant with increases in the rate of false positives and a decrease in the confidence that the compounds predicted to be stable are actually accessible. Conversely, as the threshold decreases below zero, the accuracy and false positive rate decrease together as less

and less compounds meet this stricter threshold for stability. Ultimately, the conclusions we draw from setting a stability threshold of $\Delta H_d = 0$ are not affected by alternative stability thresholds.

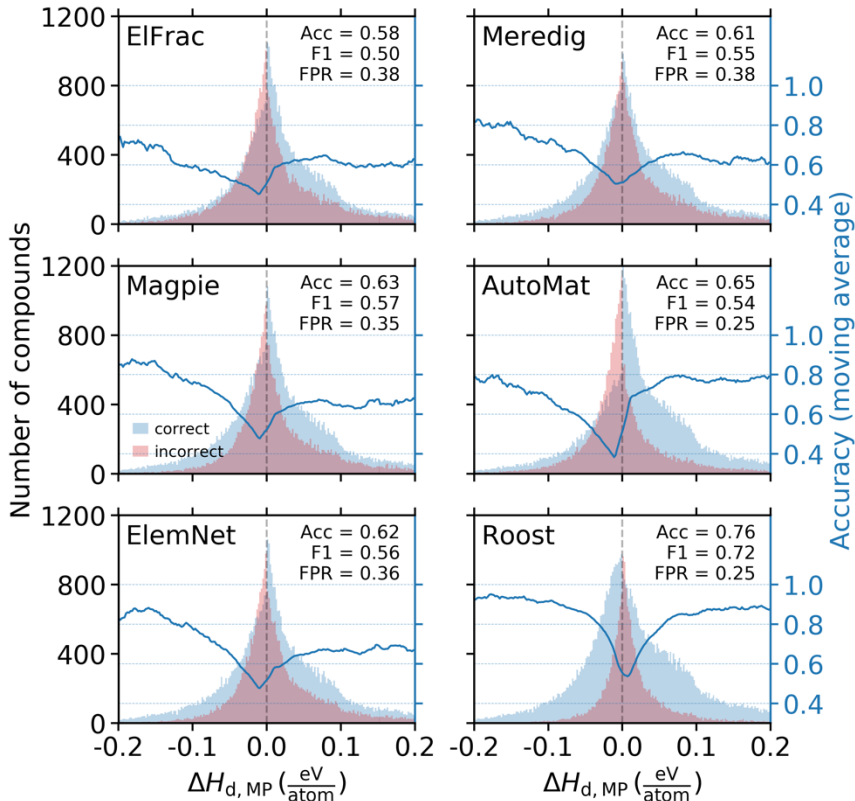


Figure 4. Classification of materials as stable ($\Delta H_d \leq 0$) or unstable ($\Delta H_d > 0$) using each of the six compositional models. “Correct” predictions are those for which the ML models and MP both predict a given material to be either stable or unstable. The histograms are binned every 5 meV/atom with respect to $\Delta H_{d,MP}$ to indicate how the correct and incorrect predictions and the number of compounds in our dataset vary as a function of the magnitude above or below the convex hull. Acc is the classification accuracy. F1 is the harmonic mean of precision and recall. FPR is the false positive rate. The moving average of the accuracy (computed within 20 meV/atom intervals) as a function of $\Delta H_{d,MP}$ is shown as a blue line (right axis). As expected, the accuracy is lowest near the chosen stability threshold of $\Delta H_{d,MP} = 0$.

Predicting stability in sparse chemical spaces

While quantifying the accuracy of ML approaches on the entire MP dataset is instructive, it does not resemble the materials discovery problem because it assesses only the limited space of compositions that have been previously explored and therefore have many stable compounds. In order to simulate a realistic materials discovery problem, we identified a set of chemical spaces within the MP dataset that are sparse in terms of stable compounds. Lithium transition metal (TM) oxides are used as the cathode material for rechargeable Li-ion batteries and have attracted

substantial attention for materials discovery in recent years. In particular, Li-Mn oxides have been considered as an alternative to LiCoO_2 utilizing less or no cobalt: e.g., spinel LiMn_2O_4 ,³⁹ layered LiMnO_2 ,⁴⁰ nickel-manganese-cobalt (NMC) cathodes,⁴¹ and disordered rock salt cathodes.⁴² For this work, the quaternary space, Li-Mn-TM-O with $\text{TM} \in \{\text{Ti}, \text{V}, \text{Cr}, \text{Fe}, \text{Co}, \text{Ni}, \text{Cu}\}$, is an attractive space to test the efficacy of these models, as it contains only 9 stable compounds and 258 more that are unstable in MP. We tested the potential for ML models to discover these stable compounds by excluding all 267 quaternary Li-Mn-TM-O compounds from the MP dataset and repeating the training of each model on ΔH_f with the remaining 84,747 compounds. We then applied each trained model to predict ΔH_f for the excluded Li-Mn-TM-O compounds and assessed their stability. Importantly, we are again concerned with DFT-calculated stability at 0 K, so we are not considering the potential for compounds in this quaternary space to be stabilized due to entropic effects (e.g., configurational disorder).

The ΔH_f parity plot for these 267 Li-Mn-TM-O compounds is shown in **Figure S3** and reveals that all models have a higher accuracy predicting ΔH_f for this subset of materials than for the entire dataset (**Figure 3**). The improved prediction of ΔH_f is likely because the compounds in this subset have strongly negative ΔH_f and are well-represented by the thousands of transition metal and lithium-containing oxides that comprise the MP dataset. Despite this improved accuracy on ΔH_f , the models all have alarmingly poor performance in predicting ΔH_d . In **Figure 5**, we show that none of the models are able to correctly detect more than three of the nine stable compounds, and even for the most successful model by this metric (*AutoMat*), the three true positives come with 24 false positives. It is noteworthy that in this experiment, the models are given a large head-start towards making these predictions because the composition space under investigation is restricted to those compounds that have DFT calculations tabulated in MP, which is biased towards stability compared to the space of all possible hypothetical compounds.

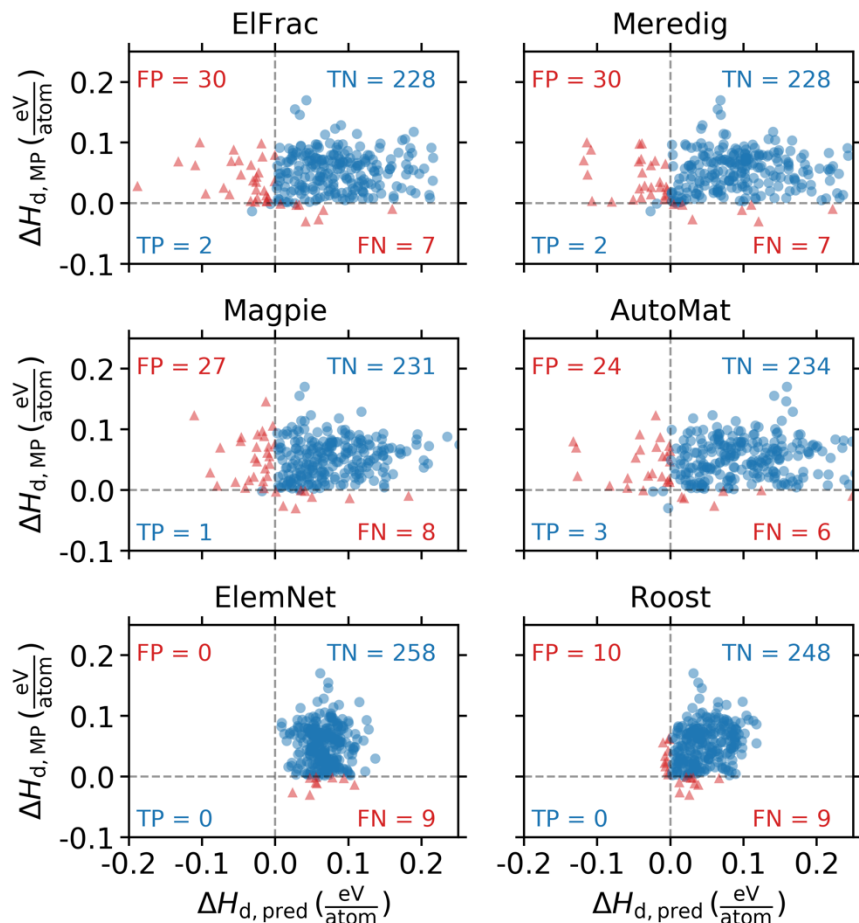


Figure 5. Re-training each model on all of MP minus 267 quaternary compounds in the Li-Mn-TM-O chemical space ($\text{TM} \in \{\text{Ti}, \text{V}, \text{Cr}, \text{Fe}, \text{Co}, \text{Ni}, \text{Cu}\}$) and obtaining ΔH_d using the predicted ΔH_f for each of the excluded compounds ($\Delta H_{d,\text{pred}}$) and comparing to stabilities available in MP, $\Delta H_{d,\text{MP}}$. FP = false positive, TP = true positive, TN = true negative, FN = false negative.

To account for the MP stability bias and more closely simulate a realistic materials discovery problem, we assessed the potential for these models to identify the nine stable MP compounds when considering a much larger composition space. Using the approach defined in Ref. ¹⁴, we produced 13,392 additional quaternary compounds in these seven Li-Mn-TM-O chemical spaces that obey simple electronegativity- and charge-based rules. For this expanded space of quaternary compounds, we used each compositional model (trained on all of MP minus the 267 Li-Mn-TM-O compounds) to predict ΔH_f and assessed their stability (**Table 1**). The compositional models each predict ~4-5% of these compounds to be stable, and all of the models fail to accurately predict the stability of more than one of the nine compounds that are actually stable in MP. A remarkable 159 compounds are predicted to be stable by all six models and 1,294 unique compounds are predicted to be stable by at least one model. While it is likely that the space

of stable quaternary compounds in the Li-Mn-TM-O space has not yet been fully explored in MP (or by extension, the ICSD), our intuition suggests it is highly unlikely that the number of new stable materials in this well-studied space is orders of magnitude larger than the number of known stable materials. The false positive rates obtained on the entire MP dataset shown in **Figure 4** suggest ~25-38% of these predicted stable Li-Mn-TM-O compounds are not actually stable, and these rates are likely underestimated, as discussed previously. The magnitude of compounds predicted to be stable by the ML models, and their false positive rates, imply that these models will inevitably identify a large number of unstable materials as candidates for further analysis (either with DFT calculations or experimental synthesis). This substantially impedes the capability of these formation energy models to accelerate the discovery of novel compounds that can be synthesized.

Table 1. Predictions in the expanded Li-Mn-TM-O (TM \in {Ti, V, Cr, Fe, Co, Ni, Cu}) composition space. Candidate compounds were generated by combining all quaternary MP compounds in this space along with quaternary compounds generated by the approach described in Ref. ¹⁴, resulting in 13,659 candidates. Among these candidates, 9 compounds are calculated to be stable in MP. The stability of all candidates was assessed using each compositional model for ΔH_f . Note that while all models correctly predict 1 of 9 MP-stable compounds to be stable, this compound is not the same for all models.

	<i>ElFrac</i>	<i>Meredig</i>	<i>Magpie</i>	<i>AutoMat</i>	<i>ElemNet</i>	<i>Roost</i>
candidate compounds	13,659	13,659	13,659	13,659	13,659	13,659
stable compounds in MP	9	9	9	9	9	9
compounds predicted stable	685	528	619	541	470	507
% predicted stable	5.0	3.9	4.5	4.0	3.4	3.7
pred. stable and stable in MP	1	1	1	1	1	1

Direct training on decomposition energy

An alternative approach to consider is to train directly on ΔH_d instead of using ML-predicted ΔH_f to obtain ΔH_d through the convex hull construction. Note that direct training on ΔH_d is complicated by the fact that ΔH_d for a given compound is dependent upon ΔH_f for other compounds within a given chemical space. This is unlike ΔH_f , which is intrinsic to a single compound. To assess the capability of each representation to directly predict stability, we repeated the analysis shown in **Figures 3-5** and **Table 1** but training each model on ΔH_d . The performance of each model on the MP Li-Mn-TM-O dataset is shown in **Figure 6**, the performance on the expanded Li-Mn-TM-O space in **Table S1**, and results for ΔH_f and ΔH_d on the entire MP dataset

are shown in **Figure S3** and **Figures S4-S5**, respectively. While the prediction accuracy (MAE and stability classification) on the entire MP dataset is typically comparable to or slightly better when training on ΔH_d (**Figures S4-S5**) instead of ΔH_f (**Figures 3-4**), the capability of the trained model to predict stability in sparse chemical spaces is even worse than when training on ΔH_f (**Figure 6, Table S1**). None of the models are able to identify even one of the nine MP-stable quaternary compounds from the set of 267 Li-Mn-TM-O compounds in MP, and every model predicts all 267 Li-Mn-TM-O compounds to be unstable (**Figure 6**). It is especially notable that for all models except *Roost* and *ElemNet*, the predictions for all 267 quaternary compounds fall in a very small window ($0.040 \text{ eV/atom} < \Delta H_{d,\text{pred}} < 0.082 \text{ eV/atom}$), suggesting the models only learn that all compounds in this space should be within the vicinity of the convex hull and do nothing to distinguish between chemically similar compounds. When the space of potential compounds is expanded to 13,659 compounds, only *Roost* and *ElemNet* predict any compound to be stable, but again, none of the nine MP stable compounds are predicted to be stable by any model (**Table S1**). Beyond the poor performance associated with these models, the direct prediction of ΔH_d is difficult to physically motivate because unlike ΔH_f , ΔH_d is not an intrinsic property of a material but depends on the energy at other compositions with which it may be in competition. This non-locality of ΔH_d also depends on the completeness of a given phase diagram: as new materials are discovered in a chemical space, ΔH_d is subject to change for any compound in that space, even if that compound's energy itself does not change, complicating the application of ML models trained on ΔH_d .

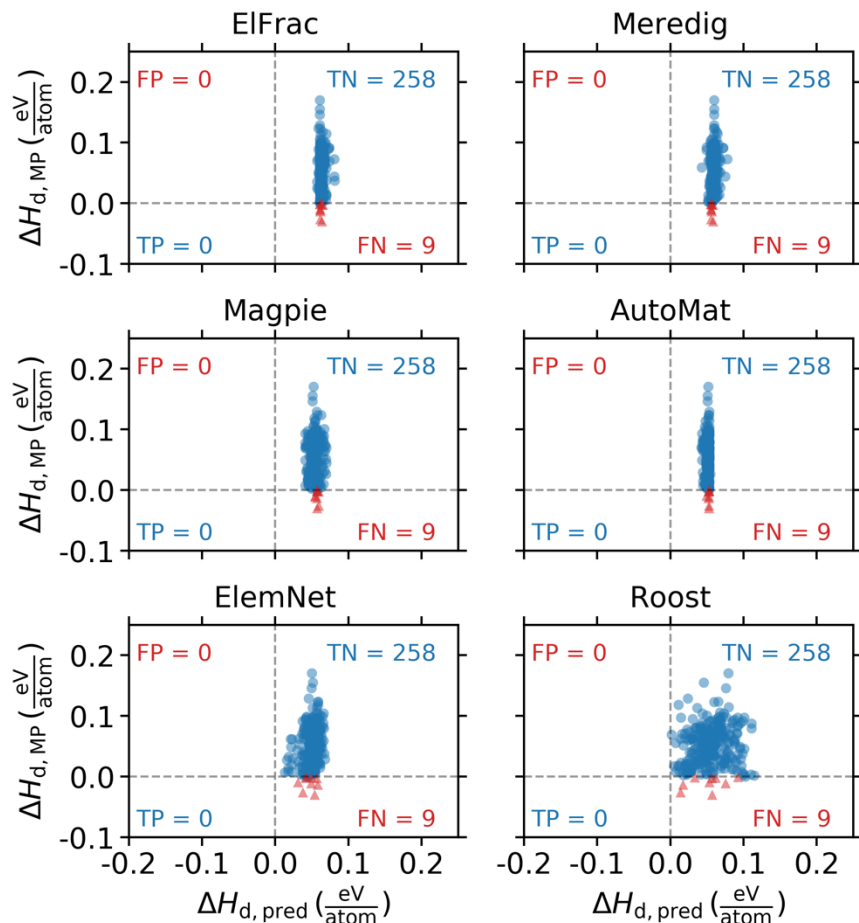


Figure 6. Re-training each model directly on ΔH_d on all of MP minus 267 quaternary compounds in the Li-Mn-TM-O chemical space (TM \in {Ti, V, Cr, Fe, Co, Ni, Cu}). Stability is determined by these direct predictions of $\Delta H_{d,\text{pred}}$ and compared to stabilities available in MP, $\Delta H_{d,\text{MP}}$. FP = false positive, TP = true positive, TN = true negative, FN = false negative.

Revisiting stability predictions with a structural representation

In addition to compositional models, representations that rely on the crystal structure for predicting formation energy have also received substantial attention in recent years.^{25,27–29,43–45} These models perform a different task than compositional models because they evaluate the property of a material given both the composition and the structure. Nevertheless, it is interesting to assess whether these structural models can predict stability with improved accuracy relative to models that are only given composition.

Here we take the crystal graph convolutional neural network (CGCNN)²⁵ as a representative example of existing structural models. CGCNN is a flexible framework that uses message passing over the atoms and bonds of a crystal (see **Methods** for training details). In

Figure 7, we show the performance of CGCNN on the same set of analyses as were shown for the compositional models in **Figures 2-5**: learning ΔH_f (**Figure 7a**), constructing convex hulls with those predicted ΔH_f to generate ΔH_d (**Figure 7b**), assessing the capability of these $\Delta H_{d,\text{pred}}$ values to classify materials as stable or unstable (**Figure 7c**), and probing the ability for this model to predict stability in the sparse Li-Mn-TM-O space (**Figure 7d**). It is clear that CGCNN improves substantially upon the direct prediction of ΔH_f (**Figure 7a**) and the implicit prediction of ΔH_d (**Figure 7b**), reducing the MAE by $\sim 50\%$ compared with the best performing compositional model (*Roost*). The extremely inaccurate predictions of ΔH_d near $\Delta H_{d,\text{pred}} = 0$ or $\Delta H_{d,\text{MP}} = 0$ that are observed in **Figure 3** for most compositional models, is also no longer present with CGCNN (**Figure 7b**). CGCNN displays an improved classification accuracy (80%) and a narrow distribution of incorrect stability predictions, only disagreeing with MP regarding the stability of compounds within the vicinity of $\Delta H_{d,\text{MP}} = 0$ (**Figure 7c**). Most impressively, CGCNN is relatively successful at finding the needles in the excluded Li-Mn-TM-O haystack, recovering five of the nine stable compounds with only six false positives (**Figure 7d**). In addition to the improved predictive accuracy, the parity plot for this excluded set looks fundamentally different than for the compositional models. In the compositional models (**Figure 5**), the parity plot is scattered, and there is effectively no linear correlation between the actual and predicted ΔH_d , whereas for CGCNN, there is a strong linear correlation (**Figure 7d**). It is not clear whether this improved correlation arises from beneficial error cancellation within each chemical space or from reducing the overall MAE from ~ 0.06 eV/atom (for the best performing compositional model – *Roost*) to ~ 0.03 eV/atom (for CGCNN).

The non-incremental improvement in stability predictions that arises from including structure in the representation is a strong endorsement for structural models and also sheds insight into the structural origins of material stability. While the thermodynamic driving force for forming a compound from its elements (formation energy) can be learned with high accuracy from only the composition, the structure dictates the subtle differences in thermodynamic driving force between chemically similar compounds and enables accurate machine learning predictions of material stability (decomposition energy). However, the glaring limitation of this approach is that it requires the structure as input, and the structure of new materials that are yet to be discovered is not known *a priori*. For example, because we do not know the ground-state structure for an arbitrary composition, we cannot repeat the test where we assess the ability of the ML model to find the

stable Li-Mn-TM-O compounds among a large set of candidate compositions. While CGCNN shows substantially improved performance in predicting material stability, these results are obtained using the DFT-optimized ground-state crystal structures as input.

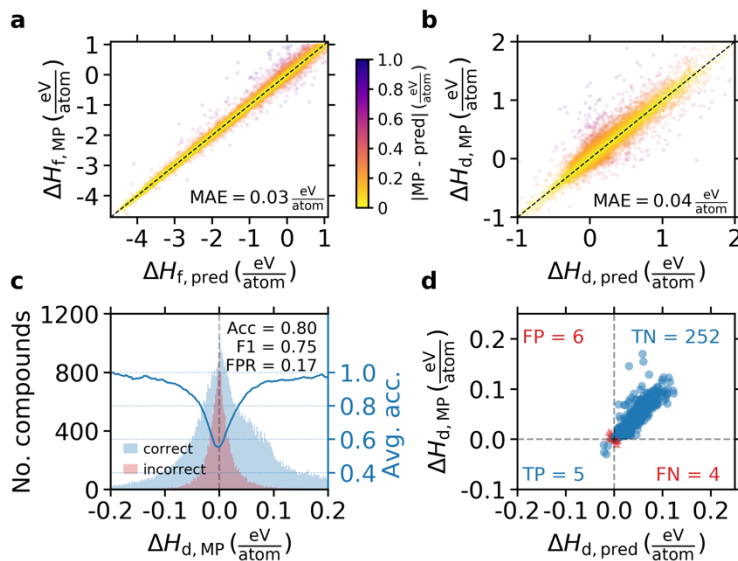


Figure 7. a) Repeating the analysis shown in **Figure 2** using CGCNN. Annotations are as in **Figure 2**. b) Repeating the analysis shown in **Figure 3** using CGCNN. Annotations are as in **Figure 3**. c) Repeating the analysis shown in **Figure 4** using CGCNN. Annotations are as in **Figure 4**. d) Repeating the analysis shown in **Figure 5** using CGCNN. Annotations are as in **Figure 5**.

Outlook

There have been a number of recent successes in the application of machine learning for materials design problems. These models have given the impression that ML can predict formation energies with near-DFT accuracy.^{20,23,46} However, the critical question of whether this implies that compound stability can be predicted by ML has not been rigorously assessed. In this work, we show that while indeed existing ML models can predict ΔH_f with relatively high accuracy from the chemical formula, they are insufficient to accurately distinguish stable from unstable compounds within an arbitrary chemical space. The error in predicting DFT-calculated ΔH_f by ML models is often compared favorably to the error DFT makes in predicting ΔH_f relative to experimentally obtained values. This comparison neglects the fact that the DFT-calculated ΔH_f differs systematically with experiment, whereas ML predictions do not. For DFT calculations, this leads to substantially lower errors for stability predictions (ΔH_d) than for ΔH_f . A similar cancellation of errors does not occur for ML models and the errors in ΔH_d are comparable to ΔH_f , inhibiting accurate predictions of material stability. Hence, while the claim that ML-predicted

formation energies have similar errors as DFT compared to experiment is technically correct, it does not imply that in their current state ML models are as useful as DFT, or that ML can replace DFT for the computationally guided discovery of novel compounds. As new ML models for formation energy are developed, it is imperative to assess their viability as inputs for stability predictions and, most critically, for problems that resemble how the models would be implemented to address emerging materials design problems. In this work, we present a set of tests that facilitate this assessment and allow for direct comparison to existing ML models. All data and code required to repeat this set of stability analyses for the models shown in this work or any new model is available at <https://github.com/CJBartel/TestStabilityML>.

Methods

Materials Project data

All entries in the Materials Project (MP)⁹ database were queried on July 26, 2019 using the Materials Project API.⁴⁷ This produced 85,014 unique non-elemental chemical formulas. For each chemical formula, we obtained the formation energy per atom, ΔH_f , for all structures having that formula, and used the most negative (ground-state) ΔH_f for training the models and obtaining ΔH_d by the convex hull construction. MP applies a correction scheme to improve the agreement between DFT-calculated thermodynamic properties (ΔH_f and ΔH_d) and experiment.^{15,48,49} Additional details on the MP calculation procedure can be found at <https://materialsproject.org/docs/calculations>.

Although the MP database contains a wide range of inorganic crystalline solids, it is an evolving resource that periodically includes more and more compounds as they are discovered or calculated by the community. As such, the calculated ΔH_d that were used for training and testing each model are subject to change over time as new stable materials are added to the database. This fact is not unique to MP and is inherent in all open materials databases that would be considered for training and evaluating machine learning models on large datasets of DFT calculations.

General training approach

Five-fold cross validation was used to produce the model-predicted $\Delta H_{f,\text{pred}}$ shown in **Figure 2**. Each predicted value corresponds with the prediction made on that compound when it was in the validation set (i.e., not used for training). $\Delta H_{f,\text{pred}}$ was then used in the convex hull

analysis to generate $\Delta H_{d,\text{pred}}$ shown in **Figure 3**, from which stability classifications were made as shown in **Figure 4**. For the Li-Mn-TM-O examples (**Figure 5** and **Table 1**), each model was trained on all MP entries except those 267 quaternary compounds belonging to the Li-Mn-TM-O chemical spaces. An analogous approach was used when training on ΔH_d instead of ΔH_f to generate the results shown in **Figure 6**, **Figure S4-S5**, and **Table S1**.

Compositional model training

Three of the compositional models—*ElFrac*, *Meredig*,²⁰ and *Magpie*²¹—were implemented using *matminer*⁵⁰ and trained using gradient boosting as implemented in XGBoost³⁰ with 200 trees and a maximum depth of 5. Preliminary tests showed XGBoost and these hyperparameters led to the highest accuracy of tested algorithms. *AutoMat*²² was used as implemented in Ref. ²². *Roost*²⁴ was trained for 500 epochs using an Adam optimizer with an initial learning rate of 5×10^{-4} and an L_1 loss function. *ElemNet* was implemented as described in Ref. ²³ using the Keras machine learning framework⁵¹. *ElemNet* was trained using an initial learning rate of 10^{-4} with an Adam optimizer for 200 epochs. 10% of the input data was set aside for validation, and the model weights from the epoch with best loss on the validation set were used for predictions. The code used to train and evaluate all models is available at <https://github.com/CJBartel/TestStabilityML>.

CGCNN training

We used a nested five-fold cross-validation to train the CGCNN²⁵ model for the MP ΔH_f dataset. As a general procedure for cross-validation, the dataset was split into five groups and each group was iteratively taken as a hold-out test set. For each fold, we split the training set to 75% training and 25% validation, thus the overall ratio of training, validation, and test was 60%, 20%, and 20%, respectively. The CGCNN model was iteratively updated by minimizing the loss (mean squared error, MSE) on the training set, and the validation score (mean absolute error, MAE) was monitored after each epoch. After 1000 epochs, the model with the best validation score was selected and then evaluated on the hold-out test set. Results of the five-fold hold-out test sets were accumulated as the final predictions of the dataset.

For the Li-Mn-TM-O case in which the test set is defined, we split the remaining compounds into five groups and iteratively took each group as the validation set (20%) and the remaining as

the training set (80%). The best CGCNN model of each fold was selected as the one with the best validation score (MAE). We then applied the five CGCNN models to the 267 Li-Mn-TM-O test compounds and used the average of the predicted ΔH_f for each model.

Acknowledgments

This work was primarily funded by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, Materials Sciences and Engineering Division under Contract No. DE-AC02-05-CH11231 (Materials Project program KC23MP). High Performance Computing resources were sponsored by the U.S. Department of Energy's Office of Energy Efficiency and Renewable Energy, located at NREL. This research also used the Savio computational cluster resource provided by the Berkeley Research Computing program at the University of California, Berkeley (supported by the UC Berkeley Chancellor, Vice Chancellor for Research, and Chief Information Officer) and the Lawrence computational cluster resource provided by the IT Division at the Lawrence Berkeley National Laboratory (Supported by the Director, Office of Science, Office of Basic Energy Sciences, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231).

Author contributions

CJB, AT, and GC conceived the project. CJB, AT, QW, and AD designed the project. AT, QW, and AD implemented the machine learning models. CJB performed the stability analysis, processed the results, and drafted the manuscript. AJ and GC supervised the project. All authors reviewed and edited the manuscript.

Data availability

A public repository at <https://github.com/CJBartel/TestStabilityML> contains the following items: code for training each compositional model in this work, code for assessing the stability of compounds given predicted formation energies (for the models shown in this work or any new model), the formation and decomposition energy data for all models studied in this work, and code for generating all figures and tables in the main text and **Supplementary Information**.

Competing interests

The authors declare no competing interests.

References

1. Himanen, L., Geurts, A., Foster, A. S. & Rinke, P. Data-Driven Materials Science: Status, Challenges, and Perspectives. *Advanced Science* **0**, 1900808 (2019).
2. Schleder, G. R., Padilha, A. C. M., Acosta, C. M., Costa, M. & Fazzio, A. From DFT to machine learning: recent approaches to materials science—a review. *Journal of Physics: Materials* **2**, 032001 (2019).
3. Goldsmith, B. R., Esterhuizen, J., Liu, J.-X., Bartel, C. J. & Sutton, C. A. Machine learning for heterogeneous catalyst design and discovery. *AIChE-Journal* **64**, 2311–2323 (2018).
4. Butler, K. T., Davies, D. W., Cartwright, H., Isayev, O. & Walsh, A. Machine learning for molecular and materials science. *Nature* **559**, 547–555 (2018).
5. Ramprasad, R., Batra, R., Pilia, G., Mannodi-Kanakkithodi, A. & Kim, C. Machine learning in materials informatics: recent applications and prospects. *npj Computational Materials* **3**, 54 (2017).
6. Schmidt, J., Marques, M. R. G., Botti, S. & Marques, M. A. L. Recent advances and applications of machine learning in solid-state materials science. *npj Computational Materials* **5**, 83 (2019).
7. Pizzi, G., Cepellotti, A., Sabatini, R., Marzari, N. & Kozinsky, B. AiiDA: automated interactive infrastructure and database for computational science. *Computational Materials Science* **111**, 218–230 (2016).
8. Kirklin, S. *et al.* The Open Quantum Materials Database (OQMD): assessing the accuracy of DFT formation energies. *npj Computational Materials* **1**, 15010 (2015).
9. Jain, A. *et al.* Commentary: The Materials Project: A materials genome approach to accelerating materials innovation. *APL Materials* **1**, 011002 (2013).
10. Draxl, C. & Scheffler, M. NOMAD: The FAIR concept for big data-driven materials science. *MRS Bulletin* **43**, 676–682 (2018).
11. Curtarolo, S. *et al.* AFLOW: An automatic framework for high-throughput materials discovery. *Computational Materials Science* **58**, 218–226 (2012).
12. Choudhary, K., Kalish, I., Beams, R. & Tavazza, F. High-throughput Identification and Characterization of Two-dimensional Materials using Density functional theory. *Scientific Reports* **7**, 5179 (2017).

13. Hellenbrandt, M. The Inorganic Crystal Structure Database (ICSD)—Present and Future. *Crystallography Reviews* **10**, 17–22 (2004).
14. Davies, D. W. *et al.* Computational Screening of All Stoichiometric Inorganic Materials. *Chem* **1**, 617–627 (2016).
15. Hautier, G., Ong, S. P., Jain, A., Moore, C. J. & Ceder, G. Accuracy of density functional theory in predicting formation energies of ternary oxides from binary oxides and its implication on phase stability. *Phys. Rev. B* **85**, 155208 (2012).
16. Bartel, C. J., Weimer, A. W., Lany, S., Musgrave, C. B. & Holder, A. M. The role of decomposition reactions in assessing first-principles predictions of solid stability. *npj Computational Materials* **5**, 4 (2019).
17. Bartel, C. J. *et al.* Physical descriptor for the Gibbs energy of inorganic crystalline solids and temperature-dependent materials chemistry. *Nature communications* **9**, 4168 (2018).
18. Ong, S. P., Wang, L., Kang, B. & Ceder, G. Li–Fe–P–O₂ Phase Diagram from First Principles Calculations. *Chem. Mater.* **20**, 1798–1807 (2008).
19. Zunger, A. Inverse design in search of materials with target functionalities. *Nature Reviews Chemistry* **2**, 0121 (2018).
20. Meredig, B. *et al.* Combinatorial screening for new materials in unconstrained composition space with machine learning. *Phys. Rev. B* **89**, 094104 (2014).
21. Ward, L., Agrawal, A., Choudhary, A. & Wolverton, C. A general-purpose machine learning framework for predicting properties of inorganic materials. *npj Computational Materials* **2**, 16028 (2016).
22. AutoMatminer. <https://github.com/hackingmaterials/automatminer>.
23. Jha, D. *et al.* ElemNet: Deep Learning the Chemistry of Materials From Only Elemental Composition. *Scientific Reports* **8**, 17593 (2018).
24. Goodall, Rhys & Lee, Alpha. Predicting materials properties without crystal structure: Deep representation learning from stoichiometry. *arXiv pre-print 1910.00617*.
25. Xie, T. & Grossman, J. C. Crystal Graph Convolutional Neural Networks for an Accurate and Interpretable Prediction of Material Properties. *Physical Review Letters* **120**, 145301 (2018).
26. Deml, A. M., O’Hayre, R., Wolverton, C. & Stevanović, V. Predicting density functional theory total energies and enthalpies of formation of metal-nonmetal compounds by linear regression. *Phys. Rev. B* **93**, 085142 (2016).

27. Faber, F. A., Lindmaa, A., von Lilienfeld, O. A. & Armiento, R. Machine Learning Energies of 2 Million Elpasolite ABC2D6 Crystals. *Physical Review Letters* **117**, 135502 (2016).
28. Ward, L. *et al.* Including crystal structure attributes in machine learning models of formation energies via Voronoi tessellations. *Phys. Rev. B* **96**, 024104 (2017).
29. Chen, C., Ye, W., Zuo, Y., Zheng, C. & Ong, S. P. Graph Networks as a Universal Machine Learning Framework for Molecules and Crystals. *Chem. Mater.* **31**, 3564–3572 (2019).
30. Chen, T. & Guestrin, C. XGBoost: A Scalable Tree Boosting System. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '16* 785–794 (2016)
doi:10.1145/2939672.2939785.
31. Stevanović, V., Lany, S., Zhang, X. & Zunger, A. Correcting density functional theory for accurate predictions of compound enthalpies of formation: Fitted elemental-phase reference energies. *Phys. Rev. B* **85**, 115104 (2012).
32. Pandey, M. & Jacobsen, K. W. Heats of formation of solids with error estimation: The mBEEF functional with and without fitted reference energies. *Phys. Rev. B* **91**, 235201 (2015).
33. Zhang, Y. *et al.* Efficient first-principles prediction of solid stability: Towards chemical accuracy. *npj Computational Materials* **4**, 9 (2018).
34. Miedema, A. R. Simple model for alloys. *Philips Tech. Rev.*, v. 33, no. 6, pp. 149-160 (1973).
35. Isaacs, E. B. & Wolverton, C. Performance of the strongly constrained and appropriately normed density functional for solid-state materials. *Phys. Rev. Materials* **2**, 063801 (2018).
36. Aykol, M., Dwaraknath, S. S., Sun, W. & Persson, K. A. Thermodynamic limit for synthesis of metastable inorganic materials. *Sci Adv* **4**, eaaq0148 (2018).
37. Sun, W. *et al.* The thermodynamic scale of inorganic crystalline metastability. *Sci Adv* **2**, e1600225 (2016).
38. Xu, K. *et al.* Show, Attend and Tell: Neural Image Caption Generation with Visual Attention. *preprint at <https://arxiv.org/abs/1502.03044>*.
39. Thackeray, M. M., David, W. I. F., Bruce, P. G. & Goodenough, J. B. Lithium insertion into manganese spinels. *Materials Research Bulletin* **18**, 461–472 (1983).
40. Armstrong, A. R. & Bruce, P. G. Synthesis of layered LiMnO₂ as an electrode for rechargeable lithium batteries. *Nature* **381**, 499–500 (1996).

41. Thackeray, M. M. *et al.* LiMnO₃-stabilized LiMO₂ (M = Mn, Ni, Co) electrodes for lithium-ion batteries. *J. Mater. Chem.* **17**, 3112–3125 (2007).
42. Lee, J. *et al.* Reversible Mn²⁺/Mn⁴⁺ double redox in lithium-excess cathode materials. *Nature* **556**, 185–190 (2018).
43. Noh, J. *et al.* Inverse Design of Solid-State Materials via a Continuous Representation. *Matter* doi:10.1016/j.matt.2019.08.017.
44. Isayev, O. *et al.* Universal fragment descriptors for predicting properties of inorganic crystals. *Nature Communications* **8**, 15679 (2017).
45. Bartók, A. P., Kondor, R. & Csányi, G. On representing chemical environments. *Phys. Rev. B* **87**, 184115 (2013).
46. Jha, D. *et al.* Enhancing materials property prediction by leveraging computational and experimental data using deep transfer learning. *Nature Communications* **10**, 5316 (2019).
47. Ong, S. P. *et al.* The Materials Application Programming Interface (API): A simple, flexible and efficient API for materials data based on REpresentational State Transfer (REST) principles. *Computational Materials Science* **97**, 209–215 (2015).
48. Jain, A. *et al.* Formation enthalpies by mixing GGA and GGA+U calculations. *Phys. Rev. B* **84**, 045115 (2011).
49. Wang, L., Maxisch, T. & Ceder, G. Oxidation energies of transition metal oxides within the GGA+U framework. *Phys. Rev. B* **73**, 195107 (2006).
50. Ward, L. *et al.* Matminer: An open source toolkit for materials data mining. *Computational Materials Science* **152**, 60–69 (2018).
51. Chollet, F. & others. Keras. (2015).

Supplementary Information

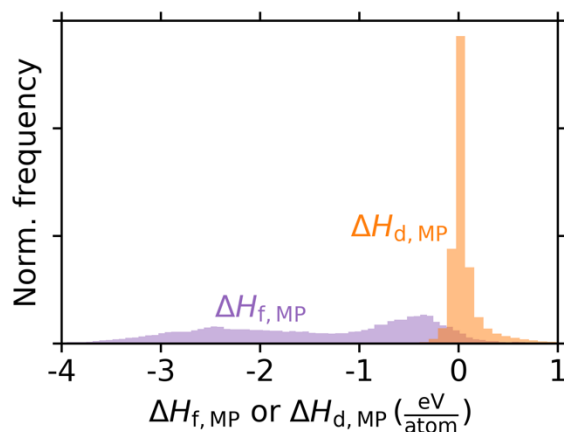


Figure S1. Distribution of ΔH_f and ΔH_d in Materials Project, with 85 bins and normalized such that the integral of the distribution is equal to 1.




Figure S2. Receiver operating characteristic (ROC) curves for each model trained on ΔH_f . TPR is the true positive rate and FPR the false positive rate. The colorbar indicates the stability threshold – i.e., a compound is classified as “stable” if ΔH_d is less than the stability threshold. Note that the models are trained on ΔH_f and are therefore insensitive to this changing threshold. Instead, the choice of threshold simply allows for an expanded analysis of the ΔH_f model performance on ΔH_d predictions.

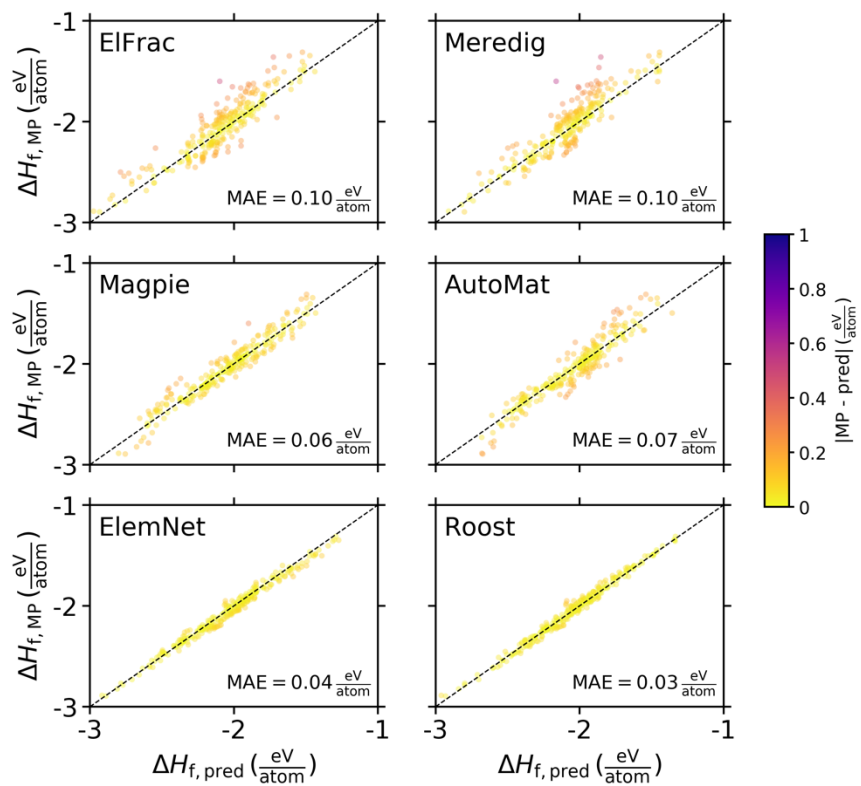


Figure S3. Re-training each model on all of MP minus 267 quaternary compounds in the Li-Mn-TM-O chemical space ($\text{TM} \in \{\text{Ti}, \text{V}, \text{Cr}, \text{Fe}, \text{Co}, \text{Ni}, \text{Cu}\}$) and predicting ΔH_f for each of the excluded compounds ($\Delta H_{f, \text{pred}}$) and comparing to MP, $\Delta H_{f, \text{MP}}$. All annotations are the same as in **Figure 2**.




Figure S4. Reproducing **Figure 3** but training on ΔH_d instead of ΔH_f . All annotations are the same as in **Figure 3**.

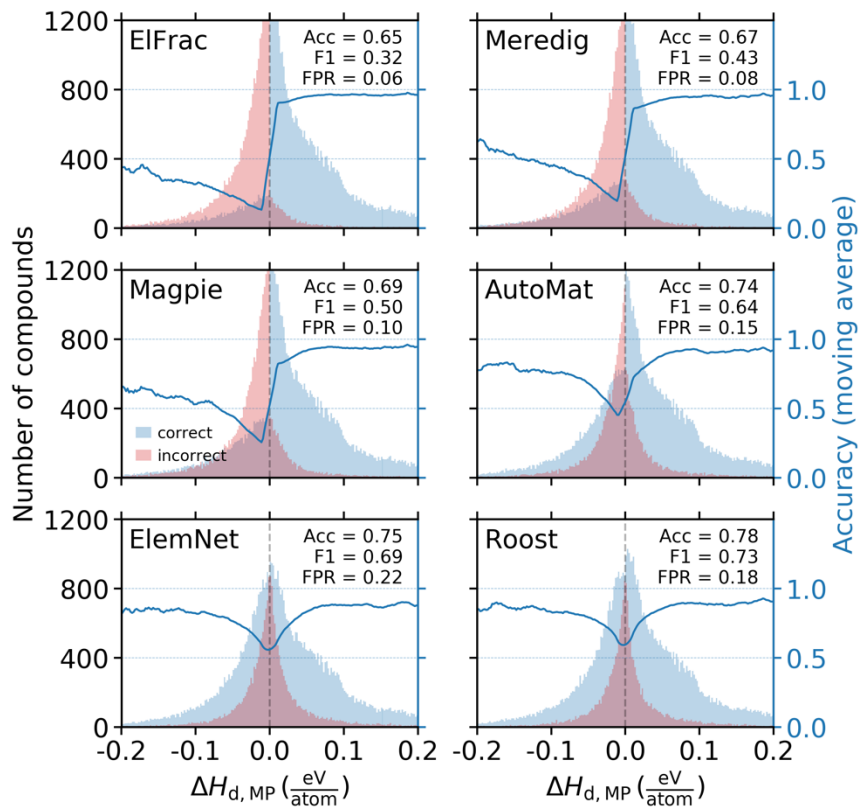


Figure S5. Reproducing **Figure 4**, but training on ΔH_d instead of ΔH_f . All annotations are the same as in **Figure 4**.

Table S1. Reproducing **Table 1** but training on ΔH_d instead of ΔH_f .

	<i>ElFrac</i>	<i>Meredig</i>	<i>Magpie</i>	<i>AutoMat</i>	<i>ElemNet</i>	<i>Roost</i>
candidate compounds	13,659	13,659	13,659	13,659	13,659	13,659
stable compounds in MP	9	9	9	9	9	9
compounds predicted stable	0	0	0	0	5	299
% predicted stable	0	0	0	0	0	2.2
pred. stable and stable in MP	0	0	0	0	0	0