

## ARTICLE OPEN

# A general-purpose machine learning framework for predicting properties of inorganic materials

Logan Ward<sup>1</sup>, Ankit Agrawal<sup>2</sup>, Alok Choudhary<sup>2</sup> and Christopher Wolverton<sup>1</sup>

A very active area of materials research is to devise methods that use machine learning to automatically extract predictive models from existing materials data. While prior examples have demonstrated successful models for some applications, many more applications exist where machine learning can make a strong impact. To enable faster development of machine-learning-based models for such applications, we have created a framework capable of being applied to a broad range of materials data. Our method works by using a chemically diverse list of attributes, which we demonstrate are suitable for describing a wide variety of properties, and a novel method for partitioning the data set into groups of similar materials to boost the predictive accuracy. In this manuscript, we demonstrate how this new method can be used to predict diverse properties of crystalline and amorphous materials, such as band gap energy and glass-forming ability.

*npj Computational Materials* (2016) **2**, 16028; doi:10.1038/npjcompumats.2016.28; published online 26 August 2016

## INTRODUCTION

Rational design of materials is the ultimate goal of modern materials science and engineering. As part of achieving that goal, there has been a large effort in the materials science community to compile extensive data sets of materials properties to provide scientists and engineers with ready access to the properties of known materials. Today, there are databases of crystal structures,<sup>1</sup> superconducting critical temperatures (<http://supercon.nims.go.jp/>), physical properties of crystalline compounds<sup>2–5</sup> and many other repositories containing useful materials data. Recently, it has been shown that these databases can also serve as resources for creating predictive models and design rules—the key tools of rational materials design.<sup>6–12</sup> These databases have grown large enough that the discovery of such design rules and models is impractical to accomplish by relying simply on human intuition and knowledge about material behaviour. Rather than relying directly on intuition, machine learning offers the promise of being able to create accurate models quickly and automatically.

To date, materials scientists have used machine learning to build predictive models for a handful of applications.<sup>13–27</sup> For example, there are now models to predict the melting temperatures of binary inorganic compounds,<sup>21</sup> the formation enthalpy of crystalline compounds,<sup>14,15,28</sup> which crystal structure is likely to form at a certain composition,<sup>5,16,29–31</sup> band gap energies of certain classes of crystals<sup>32,33</sup> and the mechanical properties of metal alloys.<sup>24,25</sup> While these models demonstrate the promise of machine learning, they only cover a small fraction of the properties used in materials design and the data sets available for creating such models. For instance, no broadly-applicable, machine-learning-based models exist for band gap energy or glass-forming ability even though large-scale databases of these properties have existed for years.<sup>2,34</sup>

Provided the large differences between the approaches used in the literature, a systematic path forward to creating accurate

machine learning models across a variety of new applications is not clear. While techniques in data analytics have advanced significantly, the development of routine methods for transforming raw materials data into the quantitative descriptions required for employing these algorithms is yet to emerge. In contrast, the chemoinformatics community benefits from a rich library of methods for describing molecular structures, which allow for standard approaches for deciding inputs into the models and, thereby, faster model development.<sup>35–37</sup> What is missing are similar flexible frameworks for building predictive models of material properties.

In this work, we present a general-purpose machine-learning-based framework for predicting the properties of materials based on their composition. In particular, we focus on the development of a set of attributes—which serve as an input to the machine learning model—that could be reused for a broad variety of materials problems. Provided a flexible set of inputs, creating a new material property model can be reduced to finding a machine learning algorithm that achieves optimal performance—a well-studied problem in data science. In addition, we employ a novel partitioning scheme to enhance the accuracy of our predictions by first partitioning data into similar groups of materials and training separate models for each group. We show that this method can be used regardless of whether the materials are amorphous or crystalline, the data are from computational or experimental studies, or the property takes continuous or discrete values. In particular, we demonstrate the versatility of our technique by using it for two distinct applications: predicting novel solar cell materials using a database of density functional theory (DFT)-predicted properties of crystalline compounds and using experimental measurements of glass-forming ability to suggest new metallic glass alloys. Our vision is that this framework could be used as a basis for quickly creating models based on the data

<sup>1</sup>Department of Materials Science and Engineering, Northwestern University, Evanston, IL, USA and <sup>2</sup>Department of Electrical Engineering and Computer Science, Northwestern University, Evanston, IL, USA.

Correspondence: C. Wolverton (c-wolverton@northwestern.edu)

Received 5 March 2016; revised 15 June 2016; accepted 10 July 2016

available in the materials databases and, thereby, initiate a major step forward in rational materials design.

## RESULTS AND DISCUSSION

The results of this study are described in two major subsections. First, we will discuss the development of our method and the characterisation of the attribute set using data from the Open Quantum Materials Database (OQMD). Next, we will demonstrate the application of this method to two distinct material problems.

### General-purpose method to create materials property models

Machine learning (ML) models for materials properties are constructed from three parts: training data, a set of attributes that describe each material, and a machine learning algorithm to map attributes to properties. For the purposes of creating a general-purpose method, we focused entirely on the attributes set because the method needs to be agnostic to the type of training data and because it is possible to utilise already-developed machine learning algorithms. Specifically, our objective is to develop a general set of attributes based on the composition that can be reused for a broad variety of problems.

The goal in designing a set of attributes is to create a quantitative representation that both uniquely defines each material in a data set and relates to the essential physics and chemistry that influence the property of interest.<sup>14,17</sup> As an example, the volume of a crystalline compound is expected to relate to the volume of the constituent elements. By including the mean volume of the constituent elements as an attribute, a machine learning algorithm could recognise the correlation between this value and the compound volume, and use it to create a predictive model. However, the mean volume of the constituent elements neither uniquely defines a composition nor perfectly describes the volumes of crystalline materials.<sup>38</sup> Consequently, one must include additional attributes to create a suitable set for this problem. Potentially, one could include factors derived from the electronegativity of the compound to reflect the idea that bond distances are shorter in ionic compounds, or the variance in atomic radius to capture the effects of polydisperse packing. The power of machine learning is that it is not necessary to know which factors actually relate to the property and how before creating a model—those relationships are discovered automatically.

The materials informatics literature is full of successful examples of attribute sets for a variety of properties.<sup>13–16,21,32,39</sup> We observed that the majority of attribute sets were primarily based on statistics of the properties of constituent elements. As an example, Meredig *et al.*<sup>15</sup> described a material based on the fraction of each element present and various intuitive factors, such as the maximum difference in electronegativity, when building models for the formation energy of ternary compounds. Ghiringhelli *et al.*<sup>14</sup> used combinations of elemental properties such as atomic number and ionisation potential to study the differences in energy between zinc-blende and rocksalt phases. We also noticed that the important attributes varied significantly depending on material property. The best attribute for describing the difference in energy between zinc-blende and rocksalt phases was found to be related to the pseudopotential radii, ionisation potential and electron affinity of the constituent elements.<sup>14</sup> In contrast, melting temperature was found to be related to atomic number, atomic mass and differences between atomic radii.<sup>21</sup> From this, we conclude that a general-purpose attribute set should contain the statistics of a wide variety of elemental properties to be adaptable.

Building on existing strategies, we created an expansive set of attributes that can be used for materials with any number of constituent elements. As we will demonstrate, this set is broad

enough to capture a sufficiently-diverse range of physical/chemical properties to be used to create accurate models for many materials problems. In total, we use a set of 145 attributes, which are described in detail and compared against other attribute sets in the Supplementary Information, that fall into four distinct categories:

1. **Stoichiometric attributes** that depend only on the fractions of elements present and not what those elements actually are. These include the number of elements present in the compound and several  $L^p$  norms of the fractions.
2. **Elemental property statistics**, which are defined as the mean, mean absolute deviation, range, minimum, maximum and mode of 22 different elemental properties. This category includes attributes such as the maximum row on periodic table, average atomic number and the range of atomic radii between all elements present in the material.
3. **Electronic structure attributes**, which are the average fraction of electrons from the *s*, *p*, *d* and *f* valence shells between all present elements. These are identical to the attributes used by Meredig *et al.*<sup>15</sup>
4. **Ionic compound attributes** that include whether it is possible to form an ionic compound assuming all elements are present in a single oxidation state, and two adaptations of the fractional 'ionic character' of a compound based on an electronegativity-based measure.<sup>40</sup>

For the third ingredient, the machine learning algorithm, we evaluate many possible methods for each individual problem. Previous studies have used machine learning algorithms including partial least-squares regression,<sup>13,29</sup> Least Absolute Shrinkage and Selection Operator (LASSO),<sup>14,33,41</sup> decision trees,<sup>15,16</sup> kernel ridge regression,<sup>17–19,42</sup> Gaussian process regression<sup>19–21,43</sup> and neural networks.<sup>22–24</sup> Each method offers different advantages, such as speed or interpretability, which must be weighed carefully for a new application. We generally approach this problem by evaluating the performance of several algorithms to find one that has both reasonable computational requirements (i.e., can be run on available hardware in a few hours) and has low error rates in cross-validation—a process that is simplified by the availability of well-documented libraries of machine learning algorithms.<sup>44,45</sup> We often find that ensembles of decision trees (e.g., rotation forests<sup>46</sup>) perform best with our attribute set. These algorithms also have the advantage of being quick to train, but are not easily interpretable by humans. Consequently, they are less suited for understanding the underlying mechanism behind a material property but, owing to their high predictive accuracy, excellent choices for the design of new materials.

We also utilise a partitioning strategy that enables a significant increase in predictive accuracy for our ML models. By grouping the data set into chemically-similar segments and training a separate model on each subset, we boost the accuracy of our predictions by reducing the breadth of physical effects that each machine learning algorithm needs to capture. For example, the physical effects underlying the stability intermetallic compounds are likely to be different than those for ceramics. In this case, one could partition the data into compounds that contain only metallic elements and another including those that do not. As we demonstrate in the examples below, partitioning the data set can significantly increase the accuracy of predicted properties. Beyond using our knowledge about the physics behind a certain problem to select a partitioning strategy, we have also explored using an automated, unsupervised-learning-based strategy for determining distinct clusters of materials.<sup>47</sup> Currently, we simply determine the partitioning strategy for each property model by searching through a large number of possible strategies and

selecting the one that minimises the error rate in cross-validation tests.

#### Justification for large attribute set

The main goal of our technique is to accelerate the creation of machine learning models by reducing or eliminating the need to develop a set of attributes for a particular problem. Our approach was to create a large attribute set, with the idea that it would contain a diverse enough library of descriptive factors such that it is likely to contain several that are well-suited for a new problem. To justify this approach, we evaluated changes in the performance of attributes for different properties and types of materials using data from the OQMD. As described in greater detail in the next section, the OQMD contains the DFT-predicted formation energy, band gap energy and volume of hundreds of thousands of crystalline compounds. The diversity and scale of the data in the OQMD make it ideal for studying changes in attribute performance using a single, uniform data set.

We found that the attributes which model a material property best can vary significantly depending on the property and type of materials in the data set. To quantify the predictive ability of each attribute, we fit a quadratic polynomial using the attribute and measured the root mean squared error of the model. We found the attributes that best describe the formation energy of crystalline compounds are based on the electronegativity of the constituent elements (e.g., maximum and mode electronegativity). In contrast, the best-performing attributes for band gap energy are the fraction of electrons in the *p* shell and the mean row in the periodic table of the constituent elements. In addition, the attributes that best describe the formation energy vary depending on the type of compounds. The formation energy of intermetallic compounds is best described by the variances in the melting temperature and number of *d* electrons between constituent elements, whereas compounds that contain at least one nonmetal are best modelled by the mean ionic character (a quantity based on electronegativity difference between constituent elements). Taken together, the changes in which attributes are the most important between these examples further support the necessity of having a large variety of attributes available in a general-purpose attribute set.

It is worth noting that the 145 attributes described in this paper should not be considered the complete set for inorganic materials. The chemical informatics community has developed thousands of attributes for predicting the properties of molecules.<sup>35</sup> What we present here is a step towards creating such a rich library of attributes for inorganic materials. While we do show in the examples considered in this work that this set of attributes is sufficient to create accurate models for two distinct properties, we expect that further research in materials informatics will add to the set presented here and be used to create models with even greater accuracy. Furthermore, many materials cannot be described simply by average composition. In these cases, we propose that the attribute set presented here can be extended with representations designed to capture additional features such as structure (e.g., Coulomb Matrix<sup>17</sup> for atomic-scale structure) or processing history. We envision that it will be possible to construct a library of general-purpose representations designed to capture structure and other characteristics of a material, which would drastically simplify the development of new machine learning models.

#### Example applications

In the following sections, we detail two distinct applications for our novel material property prediction technique to demonstrate its versatility: predicting three physically distinct properties of crystalline compounds and identifying potential metallic glass alloys. In both cases, we use the same general framework, i.e., the

same attributes and partitioning-based approach. In each case, we only needed to identify the most accurate machine learning algorithm and find an appropriate partitioning strategy. Through these examples, we discuss all aspects of creating machine-learning-based models to design a new material: assembling a training set to train the models, selecting a suitable algorithm, evaluating model accuracy and employing the model to predict new materials.

#### Accurate models for properties of crystalline compounds

DFT is a ubiquitous tool for predicting the properties of crystalline compounds, but is fundamentally limited by the amount of computational time that DFT calculations require. In the past decade, DFT has been used to generate several databases containing the  $T=0$  K energies and electronic properties of  $\sim 10^5$  crystalline compounds,<sup>2–5,48</sup> which each required millions of hours of CPU time to construct. While these databases are indisputably-useful tools, as evidenced by the many materials they have been used to design,<sup>3,49–54</sup> machine-learning-based methods offer the promise of predictions at several orders of magnitude faster rates. In this example, we explore the use of data from the DFT calculation databases as training data for machine learning models that can be used rapidly to assess many more materials than what would be feasible to evaluate using DFT.

**Training data.** We used data from the OQMD, which contains the properties of  $\sim 300,000$  crystalline compounds as calculated using DFT.<sup>2,3</sup> We selected a subset of 228,676 compounds from OQMD that represents the lowest-energy compound at each unique composition to use as a training set. As a demonstration of the utility of our method, we developed models to predict the three physically distinct properties currently available through the OQMD: band gap energy, specific volume and formation energy.

**Method.** To select an appropriate machine learning algorithm for this example, we evaluated the predictive ability of several algorithms using 10-fold cross-validation. This technique randomly splits the data set into 10 parts, and then trains a model on 9 partitions and attempts to predict the properties of the remaining set. This process is repeated using each of the 10 partitions as the test set, and the predictive ability of the model is assessed as the average performance of the model across all repetitions. As shown in Table 1, we found that creating an ensemble of reduced-error pruning decision trees using the random subspace technique had the lowest mean absolute error in cross-validation for these properties among the 10 ML algorithms we tested (of which, only 4 are listed for clarity).<sup>55</sup> Models produced using this machine learning algorithm had the lowest mean absolute error in cross-validation, and had excellent correlation coefficients of above 0.91 between the measured and predicted values for all three properties.

As a simple test for how well our band gap model can be used for discovering new materials, we simulated a search for compounds with a band gap within a desired range. To evaluate the ability of our method to locate compounds that have band gap energies within the target range, we devised a test where a model was trained on 90% of the data set and then was tasked with selecting which 30 compounds in the remaining 10% were most likely to have a band gap energy in the desired range for solar cells: 0.9–1.7 eV.<sup>56</sup> For this test, we selected a subset of the OQMD that only includes compounds that were reported to be possible to be made experimentally in the ICSD (a total of 25,085 entries) so that only band gap energy, and not stability, needed to be considered.

For this test, we compared three selection strategies for finding compounds with desirable band gap energies: randomly selecting nonmetal-containing compounds (i.e., without machine learning),

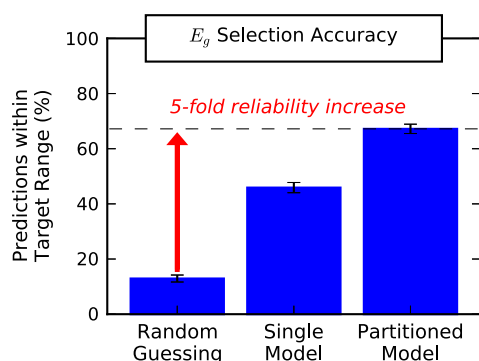


**Table 1.** Comparison of the ability of several machine learning algorithms to predict properties of materials from the OQMD

Property	Machine learning algorithm			
	Linear regression	Reduced-error pruning tree (REPTree)	Rotation forest <sup>46</sup> +REPTree	Random subspace <sup>55</sup> +REPTree
Volume (Å <sup>3</sup> per atom)	1.22	0.816	0.593	0.563
Formation energy (eV per atom)	0.259	0.126	0.0973	0.0882
Band gap energy (eV)	0.202	0.0701	0.0643	0.0645

Abbreviations: DFT, density functional theory; OQMD, Open Quantum Materials Database.

Data represents the mean absolute error in a 10-fold cross-validation test of a single model trained on the properties predicted using DFT of 228,676 crystalline compounds.



**Figure 1.** Performance of three different strategies to locate compounds with a band gap energy within a desired range: randomly selecting nonmetal-containing compounds, and two strategies using the machine-learning-based method presented in this work. The first machine learning strategy used a single model trained on the computed band gap energies of 22,667 compounds from the ICSD. The second method a model created by first partitioning the data into groups of similar materials, and training a separate model on each subset. The number of materials that were actually found to have a band gap within the desired range after 30 guesses was over 5 times larger when using our machine learning approach than when randomly selecting compounds. Error bars represent the 95% confidence interval.

using a single model trained on the entire training set to guide selection, and a model created using the partitioning approach introduced in this manuscript. As shown in Figure 1, randomly selecting a nonmetal-containing compound would result in just over 12% of the 30 selected compounds to be within the desired range of band gap energies. Using a single model trained on the entire data set, this figure dramatically improves to ~46% of selected compounds having the desired property. We found the predictive ability of our model can be increased to ~67% of predictions actually having the desired band gap energy by partitioning the data set into groups of similar compounds before training. Out of the 20 partitioning strategies we tested, we found the best composite model works by first partitioning the data set using a separate model trained to predict the expected range, but not the actual value, of the band gap energy (e.g., compounds predicted to have a band gap between 0 and 1.5 eV are grouped together), and then on whether a compound contains a halogen, chalcogen or pnictogen. Complete details of the hierarchical model are available in the Supplementary Information. By partitioning the data into smaller subsets, each of the individual machine learning models only evaluates compounds with similar chemistries (e.g., halogen-containing compounds with a band gap expected to be between 0 and 1.5 eV), which we found enhances the overall accuracy of our model.

**Table 2.** Compositions and predicted band gap energies of materials predicted using machine learning to be candidates for solar cell applications

Composition	$E_g$ (eV)
S <sub>2</sub> Hg <sub>4</sub> Cl <sub>7</sub>	1.26
V <sub>2</sub> Hg <sub>3</sub> Cl <sub>7</sub>	1.16
Mn <sub>6</sub> CCl <sub>8</sub>	1.28
Hf <sub>4</sub> S <sub>11</sub> Cl <sub>2</sub>	1.11
VCu <sub>5</sub> Cl <sub>9</sub>	1.19

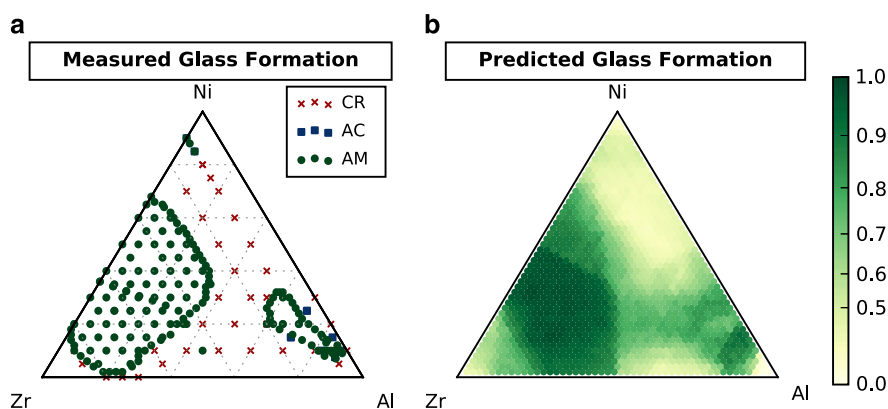
Abbreviations: DFT, density functional theory; OQMD, open quantum materials database.

Compositions represent the nominal compositions of novel ternary compounds predicted by using methods developed in ref. 15. Band gap energies were predicted using a machine learning model trained on DFT band gap energies from the OQMD<sup>2</sup> using methods described in this work.

Once we established the reliability of our model, we used it to search for new compounds (i.e., those not yet in the OQMD) with a band gap energy within the desired range for solar cells: 0.9–1.7 eV. To gain the greatest predictive accuracy, we trained our band gap model on the entire OQMD data set. Then, we used this model to predict the band gap energy of compositions that were predicted by Meredig *et al.*<sup>15</sup> to be as-yet-undiscovered ternary compounds. Out of this list of 4,500 predicted compounds, we found that 223 are likely to have favourable band gap energies. A subset with the best stability criterion (as reported in ref. 15) and band gap energy closest to 1.3 eV are shown in Table 2. As demonstrated in this example and a recent work by Sparks *et al.*,<sup>57</sup> having access to several machine learning models for different properties can make it possible to rapidly screen materials based on many design criteria. Provided the wide range of applicability of the machine learning technique demonstrated in this work and the growing availability of material property data, it may soon be possible to screen for materials based on even more properties than those considered here using models constructed based on several different data sets.

#### Locating novel metallic glass alloys

Metallic glasses possess a wide range of unique properties, such as high-wear resistance and soft magnetic behaviour, but are only possible to create at special compositions that are difficult to determine *a priori*.<sup>58</sup> The metallic glass community commonly relies on empirical rules (e.g., systems that contain many elements of different sizes are more likely to form glasses<sup>59</sup>) and extensive experimentation to locate these special compositions.<sup>55</sup> While searches based on empirical rules have certainly been successful (as evidenced by the large variety of known alloys,<sup>60</sup>) this conventional method is known to be slow and resource-intensive.<sup>61</sup> Here, we show how machine learning could be used



**Figure 2.** (a) Experimental measurements of metallic glass-forming ability in the Al–Ni–Zr ternary, as reported in ref. 34. Green circles (AM) mark compositions at which it is possible to create a fully-amorphous ribbon via melt spinning, blue squares (AC) mark compositions at which only a partially-amorphous ribbon can be formed, and red crosses (CR) mark compositions where it is not possible to form any appreciable amount of amorphous phase. (b) Predicted glass-forming ability from our machine learning model. Points are coloured based on relative likelihood of glass formation, where 1 is the mostly likely and 0 is the least. The model used to make these predictions was developed using the methods outlined in this work, and was not trained on any measurements from the Al–Ni–Zr ternary or any of its constituent binaries.

to accelerate the discovery of new alloys by using known experimental data sets to construct predictive models of glass-forming ability.

**Data.** We used experimental measurements taken from ‘Nonequilibrium Phase Diagrams of Ternary Amorphous Alloys,’ a volume of the Landolt–Börnstein collection.<sup>32</sup> This data set contains measurements of whether it is possible to form a glass using a variety of experimental techniques at thousands of compositions from hundreds of ternary phase diagrams. For our purposes, we selected 5,369 unique compositions where the ability to form an amorphous ribbon was assessed using melt spinning. In the event that multiple measurements for glass-forming ability were taken at a single composition, we assume that it is possible to form a metallic glass if at least one measurement found it was possible to form a completely amorphous sample. After the described screening steps, 70.8% of the entries in the training data set correspond to metallic glasses.

**Method.** We used the same set of 145 attributes as in the band gap example and ensembles of Random Forest classifiers<sup>62</sup> created using the random subspace technique as the machine learning algorithm, which we found to be the most accurate algorithm for this problem. This model classifies the data into two categories (i.e., can and cannot form a metallic glass) and computes the relative likelihood that a new entry would be part of each category. For the purposes of validating the model, we assume any composition predicted to have a > 50% probability of glass formation to be a positive prediction of glass-forming ability. Using a single model trained on the entire data set, we were able to create a model with 90% accuracy in 10-fold cross-validation.

As a test of the ability of our method to predict new alloys, we removed all entries that contained exclusively Al, Ni and Zr (i.e., all Al–Ni–Zr ternary compounds, and any binary formed by any two of those elements) from our training data set and then predicted the probability of an alloy being able to be formed into the amorphous state for the Al–Ni–Zr ternary system. As shown in Figure 2a, it is possible to form amorphous ribbons with melt spinning in one region along the Ni–Zr binary and in a second, Al-rich ternary region. Our model is able to accurately predict both the existence of these regions and their relative locations (see Figure 2b), which shows that models created using our

method could serve to accurately locate favourable compositions in yet-unassessed alloy systems.

We further validated the ability of our models to extrapolate to alloy systems not included in the training set by iteratively using each binary system as a test set. This procedure works by excluding all alloys that contain both of the elements in the binary, training a model on the remaining entries and then predicting the glass-forming ability of the alloys that were removed. For example, if the Al–Ni binary were being used as a test set, then  $\text{Al}_{50}\text{Ni}_{50}$  and  $\text{Al}_{50}\text{Ni}_{25}\text{Fe}_{25}$  would be removed but  $\text{Al}_{50}\text{Fe}_{50}$  and  $\text{Al}_{50}\text{Fe}_{25}\text{Zr}_{25}$  would not. This process is then repeated for all 380 unique binaries in the data set. We measured that our model has an 80.2% classification accuracy over 15,318 test entries where 71% of entries were measured to be glasses—in contrast to the 90.1% measured in 10-fold cross-validation with a similar fraction of glasses in the test set. We also found that by training separate models for alloys that contain only metallic elements and those that contain a nonmetal/metalloid it is possible to slightly increase the prediction accuracy to 80.7%—a much smaller gain than that observed in the band gap example (23%). Overall, this exclusion test strongly establishes that our model is able to predict the glass-forming ability in alloy systems that are completely unassessed.

To search for new candidate metallic glasses, we used our model to predict the probability of glass formation for all possible ternary alloys created at 2 at% spacing by any combination of elements found in the training set. Considering that the data set included 51 elements, this space includes ~24 million candidate alloys, which required ~6 h to evaluate on eight 2.2 GHz processors. To remove known alloys from our prediction results, we first removed all entries where the  $L_1$  distance between the composition vector (i.e.,  $\langle x_{\text{H}}, x_{\text{He}}, x_{\text{Li}}, \dots \rangle$ ) of the alloy and any amorphous alloy in the training set was < 30 at%. We then found the alloys with the highest predicted probability of glass formation in each binary and ternary. Eight alloys with the highest probability of glass formation are shown in Table 3. One top candidate,  $\text{Zr}_{0.38}\text{Co}_{0.24}\text{Cu}_{0.38}$ , is particularly promising considering the existence of Zr-lean Zr–Co and Zr–Cu binary alloys and Zr–Al–Co–Cu bulk metallic glasses.<sup>63</sup> To make the ability to find new metallic glasses openly available to the materials science community, we have included all of the software and data necessary to use this model in the Supplementary Information and created an interactive, web-based tool ([http://oqmd.org/static/analytics/glass\\_search.html](http://oqmd.org/static/analytics/glass_search.html)).

**Table 3.** Compositions of candidate metallic glass alloys predicted using a machine learning model trained on experimental measurements of glass-forming ability

Alloy composition	
Zr <sub>0.38</sub> Co <sub>0.24</sub> Cu <sub>0.38</sub>	Hf <sub>0.7</sub> Si <sub>0.16</sub> Ni <sub>0.14</sub>
V <sub>0.16</sub> Ni <sub>0.64</sub> B <sub>0.2</sub>	Hf <sub>0.48</sub> Zr <sub>0.16</sub> Ni <sub>0.36</sub>
Zr <sub>0.46</sub> Cr <sub>0.36</sub> Ni <sub>0.18</sub>	Zr <sub>0.48</sub> Fe <sub>0.46</sub> Ni <sub>0.06</sub>
Zr <sub>0.5</sub> Fe <sub>0.38</sub> W <sub>0.12</sub>	Sm <sub>0.22</sub> Fe <sub>0.54</sub> B <sub>0.24</sub>

These alloys were predicted to have the highest probability being able to be formed into an amorphous ribbon via melting spinning out of 24 million candidates.

## CONCLUSIONS

In this work, we introduced a general-purpose machine learning framework for predicting the properties of a wide variety of materials and demonstrated its broad applicability via illustration of two distinct materials problems: discovering new potential crystalline compounds for photovoltaic applications and identifying candidate metallic glass alloys. Our method works by using machine learning to generate models that predict the properties of a material as a function of a wide variety of attributes designed to approximate chemical effects. The accuracy of our models is further enhanced by partitioning the data set into groups of similar materials. In this manuscript, we show that this technique is capable of creating accurate models for properties as different as the electronic properties of crystalline compounds and glass formability of metallic alloys. Creating new models with our strategy requires only finding which machine learning algorithm maximises accuracy and testing different partitioning strategies, which are processes that could be eventually automated.<sup>64</sup> We envision that the versatility of this method will make it useful for a large range of problems, and help enable the quicker deployment and wider-scale use machine learning in the design of new materials.

## MATERIALS AND METHODS

All machine learning models were created using the Weka machine learning library.<sup>44</sup> The Materials Agnostic Platform for Informatics and Exploration (Magpie) was used to compute attributes, perform the validation experiments and run searches for new materials. Both Weka and Magpie are available under open-source licenses. The software, training data sets and input files used in this work are provided in the Supplementary Information associated with this manuscript.

## ACKNOWLEDGEMENTS

This work was performed under the following financial assistance award 70NANB14H012 from U.S. Department of Commerce, National Institute of Standards and Technology as part of the Center for Hierarchical Materials Design (CHiMaD). In addition, AA and AC were supported in part by the following grants: DARPA SIMPLEX award N66001-15-C-4036; NSF awards IIS-1343639 and CCF-1409601; DOE award DESC0007456; and AFOSR award FA9550-12-1-0458. LW was partially supported by the Department of Defense (DoD) through the National Defense Science & Engineering Graduate Fellowship (NDSEG) Program.

## CONTRIBUTIONS

CW conceived the project, and jointly developed the method with LW, AA and AC. LW wrote all software and performed the necessary calculations with help and guidance from AA and AC. LW led the manuscript writing, with contributions from all other authors.

## COMPETING INTERESTS

The authors declare no conflict of interest.

## REFERENCES

- Belsky, A., Hellenbrandt, M., Karen, V. L. & Luksch, P. New developments in the Inorganic Crystal Structure Database (ICSD): accessibility in support of materials research and design. *Acta Crystallogr. Sect. B Struct. Sci.* **58**, 364–369 (2002).
- Kirklin, S. et al. The Open Quantum Materials Database (OQMD): assessing the accuracy of DFT formation energies. *NPJ Comput. Mater.* **1**, 15010 (2015).
- Saal, J. E., Kirklin, S., Aykol, M., Meredig, B. & Wolverton, C. Materials design and discovery with high-throughput density functional theory: the open quantum materials database (OQMD). *JOM* **65**, 1501–1509 (2013).
- Jain, A. et al. Commentary: the materials project: a materials genome approach to accelerating materials innovation. *APL Mater.* **1**, 011002 (2013).
- Curtarolo, S. et al. AFLOWLIB.ORG: a distributed materials properties repository from high-throughput ab initio calculations. *Comput. Mater. Sci.* **58**, 227–235 (2012).
- Kalidindi, S. R. & De Graef, M. Materials data science: current status and future outlook. *Annu. Rev. Mater. Res.* **45**, 171–193 (2015).
- Kalinin, S. V., Sumpster, B. G. & Archibald, R. K. Big-deep-smart data in imaging for guiding materials design. *Nat. Mater.* **14**, 973–980 (2015).
- Rajan, K. Materials informatics: the materials ‘gene’ and big data. *Annu. Rev. Mater. Res.* **45**, 153–169 (2015).
- Rajan, K. Materials informatics. *Mater. Today* **8**, 38–45 (2005).
- Lookman, T., Alexander, F. J. & Bishop, A. R. Perspective: codesign for materials science: an optimal learning approach. *APL Mater.* **4**, 053501 (2016).
- Mulholland, G. J. & Paradiso, S. P. Perspective: Materials informatics across the product lifecycle: Selection, manufacturing, and certification. *APL Mater.* **4**, 053207 (2016).
- Agrawal, A. & Choudhary, A. Perspective: Materials informatics and big data: Realization of the ‘fourth paradigm’ of science in materials science. *APL Mater.* **4**, 053208 (2016).
- Srinivasan, S. & Rajan, K. ‘Property phase diagrams’ for compound semiconductors through data mining. *Materials (Basel)* **6**, 279–290 (2013).
- Ghiringhelli, L. M., Vybiral, J., Levchenko, S. V., Draxl, C. & Scheffler, M. Big data of materials science: critical role of the descriptor. *Phys. Rev. Lett.* **114**, 105503 (2015).
- Meredig, B. et al. Combinatorial screening for new materials in unconstrained composition space with machine learning. *Phys. Rev. B* **89**, 094104 (2014).
- Kong, C. S. et al. Information-theoretic approach for the discovery of design rules for crystal chemistry. *J. Chem. Inf. Model.* **52**, 1812–1820 (2012).
- Faber, F., Lindmaa, A., von Lilienfeld, O. A. & Armiento, R. Crystal structure representations for machine learning models of formation energies. *Int. J. Quantum Chem.* **115**, 1094–1101 (2015).
- Schütt, K. T. et al. How to represent crystal structures for machine learning: towards fast prediction of electronic properties. *Phys. Rev. B* **89**, 205118 (2014).
- Pilania, G., Wang, C., Jiang, X., Rajasekaran, S. & Ramprasad, R. Accelerating materials property predictions using machine learning. *Sci. Rep.* **3**, 2810 (2013).
- Bartók, A. P., Payne, M. C., Kondor, R. & Csányi, G. Gaussian approximation potentials: the accuracy of quantum mechanics, without the electrons. *Phys. Rev. Lett.* **104**, 136403 (2010).
- Seko, A., Maekawa, T., Tsuda, K. & Tanaka, I. Machine learning with systematic density-functional theory calculations: application to melting temperatures of single- and binary-component solids. *Phys. Rev. B* **89**, 054303 (2014).
- Hou, Z.-Y., Dai, Q., Wu, X.-Q. & Chen, G.-T. Artificial neural network aided design of catalyst for propane ammoxidation. *Appl. Catal. A Gen.* **161**, 183–190 (1997).
- Sumpster, B. & Noid, D. On the design, analysis, and characterisation of materials using computational neural networks. *Annu. Rev. Mater. Sci.* **26**, 223–277 (1996).
- Bhadeshia, H. K. D. H., Dimitriu, R. C., Forsik, S., Pak, J. H. & Ryu, J. H. Performance of neural networks in materials science. *Mater. Sci. Technol.* **25**, 504–510 (2009).
- Chatterjee, S., Muruganath, M. & Bhadeshia, H. K. D. H.  $\delta$  TRIP steel. *Mater. Sci. Technol.* **23**, 819–827 (2007).
- Hautier, G. in *Prediction and Calculation of Crystal Structures*. (eds Atahan-Evrenk, S. & Aspuru-Guzik, A.) 139–179 (Springer International Publishing, 2014).
- Yang, L. & Ceder, G. Data-mined similarity function between material compositions. *Phys. Rev. B* **88**, 224107 (2013).
- Deml, A. M., Hayre, R. O., Wolverton, C. & Stevanovic, V. Predicting density functional theory total energies and enthalpies of formation of metal-nonmetal compounds by linear regression. *Phys. Rev. B* **93**, 085142 (2016).
- Curtarolo, S., Morgan, D., Persson, K., Rodgers, J. & Ceder, G. Predicting crystal structures with data mining of quantum calculations. *Phys. Rev. Lett.* **91**, 135503 (2003).

30. Fischer, C. C., Tibbetts, K. J., Morgan, D. & Ceder, G. Predicting crystal structure by merging data mining with quantum mechanics. *Nat. Mater.* **5**, 641–646 (2006).
31. Hautier, G., Fischer, C., Ehrlacher, V., Jain, A. & Ceder, G. Data mined ionic substitutions for the discovery of new compounds. *Inorg. Chem.* **50**, 656–663 (2011).
32. Dey, P. *et al.* Informatics-aided bandgap engineering for solar materials. *Comput. Mater. Sci.* **83**, 185–195 (2014).
33. Pilania, G. *et al.* Machine learning bandgaps of double perovskites. *Sci. Rep.* **6**, 19375 (2016).
34. Kawazoe, Y., Yu, J. Z., Tsai, A. P. & Masumoto T (eds). Nonequilibrium Phase Diagrams of Ternary Amorphous Alloys. (Springer-Verlag, Berlin, Germany, 1997).
35. Todeschini, R. & Consonni, V. *Handbook of Molecular Descriptors* (Wiley-VCH Verlag GmbH, 2000).
36. Ruiz-Blanco, Y. B., Paz, W., Green, J. & Marrero-Ponce, Y. ProtDCal: a program to compute general-purpose-numerical descriptors for sequences and 3D-structures of proteins. *BMC Bioinformatics* **16**, 162 (2015).
37. Mauri, A., Consonni, V., Pavan, M. & Todeschini, R. Dragon software: an easy approach to molecular descriptor calculations. *Match Commun. Math. Comput. Chem.* **56**, 237–248 (2006).
38. Denton, A. R. & Ashcroft, N. W. Vegards law. *Phys. Rev. A* **43**, 3161–3164 (1991).
39. Villars, P., Cenzual, K., Daams, J., Chen, Y. & Iwata, S. Data-driven atomic environment prediction for binaries using the Mendeleev number. *J. Alloys Compd.* **367**, 167–175 (2004).
40. Callister, W. D. *Materials Science and Engineering: An Introduction* (Wiley, 2007).
41. Seko, A., Takahashi, A. & Tanaka, I. Sparse representation for a potential energy surface. *Phys. Rev. B* **90**, 024101 (2014).
42. Rupp, M., Tkatchenko, A., Müller, K.-R., Lilienfeld, V. & Anatole, O. Fast and accurate modeling of molecular atomization energies with machine learning. *Phys. Rev. Lett.* **108**, 58301 (2012).
43. Pyzer-Knapp, E. O., Simm, G. N. & Aspuru-Guzik, A. A Bayesian approach to calibrating high-throughput virtual screening results and application to organic photovoltaic materials. *J. Mater. Chem.* **2**, 303 (2015).
44. Hall, M. *et al.* The WEKA data mining software. *ACM SIGKDD Explor. Newslett.* **11**, 10 (2009).
45. King, D. Dlib-ml: a machine learning toolkit. *J. Mach. Learn. Res.* **10**, 1755–1758 (2009).
46. Rodríguez, J. J., Kuncheva, L. I. & Alonso, C. J. Rotation forest: a new classifier ensemble method. *IEEE Trans. Pattern Anal. Mach. Intell.* **28**, 1619–1630 (2006).
47. Meredig, B. & Wolverton, C. Dissolving the periodic table in cubic zirconia: data mining to discover chemical trends. *Chem. Mater.* **26**, 1985–1991 (2014).
48. Jain, A. *et al.* A high-throughput infrastructure for density functional theory calculations. *Comput. Mater. Sci.* **50**, 2295–2310 (2011).
49. Curtarolo, S. *et al.* The high-throughput highway to computational materials design. *Nat. Mater.* **12**, 191–201 (2013).
50. Kirklin, S., Meredig, B. & Wolverton, C. High-throughput computational screening of new Li-ion battery anode materials. *Adv. Energy Mater.* **3**, 252–262 (2013).
51. Gautier, R. *et al.* Prediction and accelerated laboratory discovery of previously unknown 18-electron ABX compounds. *Nat. Chem.* **7**, 308–316 (2015).
52. Chen, H. *et al.* Carbonophosphates: a new family of cathode materials for Li-ion batteries identified computationally. *Chem. Mater.* **24**, 2009–2016 (2012).
53. Liu, M. *et al.* Spinel compounds as multivalent battery cathodes: a systematic evaluation based on *ab initio* calculations. *Energy Environ. Sci.* **8**, 964–974 (2014).
54. Yang, K., Setyawan, W., Wang, S., Buongiorno Nardelli, M. & Curtarolo, S. A search model for topological insulators with high-throughput robustness descriptors. *Nat. Mater.* **11**, 614–619 (2012).
55. Ho, T. K. The random subspace method for constructing decision forests. *IEEE Trans. Pattern Anal. Mach. Intell.* **20**, 832–844 (1998).
56. Shockley, W. & Queisser, H. J. Detailed balance limit of efficiency of p-n junction solar cells. *J. Appl. Phys.* **32**, 510 (1961).
57. Sparks, T. D., Gaultois, M. W., Oliynyk, A., Brgoch, J. & Meredig, B. Data mining our way to the next generation of thermoelectrics. *Scr. Mater.* **111**, 10–15 (2015).
58. Wang, W. H., Dong, C. & Shek, C. H. Bulk metallic glasses. *Mater. Sci. Eng. R Rep.* **44**, 45–89 (2004).
59. Inoue, A. Stabilization of metallic supercooled liquid and bulk amorphous alloys. *Acta Mater.* **48**, 279–306 (2000).
60. Löffler, J. F. Formation of bulk metallic glasses and their composites. *MRS Bull.* **32**, 624–628 (2007).
61. Ding, S. *et al.* Combinatorial development of bulk metallic glasses. *Nat. Mater.* **13**, 494–500 (2014).
62. Breiman, L. Random forests. *Mach. Learn.* **45**, 5–32 (2001).
63. Wada, T., Zhang, T. & Inoue, A. Formation and high mechanical strength of bulk glassy alloys in Zr-Al-Co-Cu system. *Mater. Trans.* **44**, 1839–1844 (2003).
64. Thornton, C., Hutter, F., Hoos, H. H. & Leyton-Brown, K. in Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. (ACM, New York, NY, 2013).



This work is licensed under a Creative Commons Attribution 4.0 International License. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in the credit line; if the material is not included under the Creative Commons license, users will need to obtain permission from the license holder to reproduce the material. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>

© The Author(s) 2016

Supplementary Information accompanies the paper on the *npj Computational Materials* website (<http://www.nature.com/npjcompumats>)