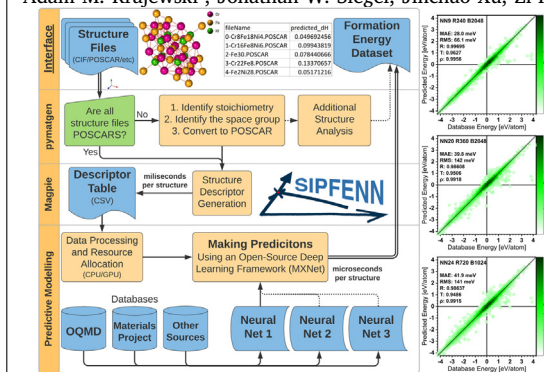


Graphical Abstract

Extensible Structure-Informed Prediction of Formation Energy with improved accuracy and usability employing neural networks

Computational Materials Science xxx (xxxx) xxx

Adam M. Krajewski^{*}, Jonathan W. Siegel, Jinchao Xu, Zi-Kui Liu



Graphical abstract and Research highlights will be displayed in online search result lists, the online contents list and the online article, but **will not appear in the article PDF file or print** unless it is mentioned in the journal specific style requirement. They are displayed in the proof pdf for review purpose only.



Full length article

Extensible Structure-Informed Prediction of Formation Energy with improved accuracy and usability employing neural networks

Adam M. Krajewski^{a,*}, Jonathan W. Siegel^b, Jinchao Xu^b, Zi-Kui Liu^a^a Department of Materials Science and Engineering, The Pennsylvania State University, USA^b Department of Mathematics, The Pennsylvania State University, USA

ARTICLE INFO

Keywords:

Machine learning
Structure-informed
Formation energy
SIPFENN

ABSTRACT

In the present paper, we introduce a new neural network-based tool for the prediction of formation energies of atomic structures based on elemental and structural features of Voronoi-tessellated materials. We provide a concise overview of the connection between the machine learning and the true material–property relationship, how to improve the generalization accuracy by reducing overfitting, how new data can be incorporated into the model to tune it to a specific material system, and preliminary results on using models to preform local structure relaxations.

The present work resulted in three final models optimized for (1) highest test accuracy on the Open Quantum Materials Database (OQMD), (2) performance in the discovery of new materials, and (3) performance at a low computational cost. On a test set of 21,800 compounds randomly selected from OQMD, they achieve a mean absolute error (MAE) of 28, 40, and 42 meV/atom, respectively. The second model provides better predictions in a test case of interest not present in the OQMD, while the third reduces the computational cost by a factor of 8.

We collect our results in a new open-source tool called SIPFENN (Structure-Informed Prediction of Formation Energy using Neural Networks). SIPFENN not only improves the accuracy beyond existing models but also ships in a ready-to-use form with pre-trained neural networks and a GUI interface. By virtue of this, it can be included in DFT calculations routines at nearly no cost.

1. Introduction

In recent years the field of material data informatics has been growing in importance thanks to the proliferation of open-access databases [1–7] and new methods being implemented to predict a wide variety of material properties [8–17]. Within these methods, machine learning (ML) and, more broadly, artificial intelligence (AI) is becoming dominant, as noted in two recent reviews [18,19], which listed a total of around 100 recent studies that attempted to solve material science problems using ML and AI techniques. These studies report benefits such as a 30-fold increase in material discovery rate when guided by an ML-model [15], or the ability to create new state-of-the-art materials in highly complex design spaces like 6-component alloys [16]. They also dive into new paradigms of materials science by handling previously unthinkable amounts of data, allowing the creation and analysis of an energy convex-hull calculated for all elements [20,21], or a concurrent analysis of all available literature texts to find paths for material synthesis [22]. In addition, some studies promise to solve significant industrial challenges such as detection of additive manufacturing flaws

with relatively simple and accessible data, but above-human pattern recognition quality and speed [17].

A common approach is to focus on the discovery of candidate materials promising a new state-of-the-art performance, which must then be validated by experiment. The mismatch between the predictions and experiment measures the quality of the model, and reducing this gap is a major challenge due to the newly designed materials often being far from known materials, combined with attention placed on regions with extraordinary predictions. However, even if design models were perfectly accurate, many predicted materials cannot be physically made in the lab. An increasing number of studies attempt to solve this challenge by focusing not only on predicting how the material will perform but also on whether it can be manufactured [23]. Generally, these include predicting materials' stability [21,24–27] and synthesizability [20,22,28] with the stability being the more constraining parameter, as it determines whether the material could be stable or metastable in the use conditions, and therefore whether it can be synthesizable. Thus,

^{*} Corresponding author.E-mail address: ak@psu.edu (A.M. Krajewski).

predicting stability through prediction of fundamental thermodynamic properties such as formation energy is of special importance.

In the present work, new ML models and a tool to quickly use them are developed to improve the process of materials discovery by efficient prediction the formation energy and streamlined incorporation into materials discovery frameworks that aim to screen billions rather than hundreds of candidates available with cost-intensive calculations like first-principles calculations based on the density functional theory (DFT).

In simple terms, every ML model is composed of three essential elements: a database, a descriptor, and an ML technique (also known as ML algorithm). The first element, databases, contain prior knowledge and are becoming increasingly shared between many studies, thanks to being open-access and often containing orders of magnitude more experimental or computational data than could be feasibly collected for a single study [1–7]. Databases used within the present paper are detailed in Section 2.4.

The second element of an ML model is the descriptor (i.e., feature vector describing the material) which determines a representation of knowledge (data from the database) in a way relevant to the problem. It is typically built from many features, also known as attributes or vector components, which usually are determined through domain knowledge to be relevant or selected through correlation analysis. All combined, these features are a representation of some state whose meaning will be problem-specific.

When treating materials on the single atomic configuration level, descriptors can be generally divided into composition-based (also known as stoichiometric, structure-invariant, or elemental) [9,29,30] and structure-informed [31–33]. The first type usually provides a more compact representation at a much lower computational cost, as calculating a composition-based descriptor often needs to involve only simple linear algebra operations such as matrix multiplication [30], or prior-knowledge-incorporating attention-based analysis of a graph representation of the composition [34]. In cases where deep neural networks (DNNs) are employed, descriptor calculation can be skipped altogether by passing a composition vector directly [29].

It is important to recognize that the descriptor choice impacts both the performance and applicability of the model. In the case of prediction of material properties, such as formation energy, selecting a composition-based descriptor, no matter how complex, limits the model to either a specific arrangement of atoms, such as BCC or amorphous, or some defined pattern of structures, such as the convex hull of lowest-energy structures. Such limitation of the problem domain, given a comparable amount of data, allows to quickly achieve much lower prediction error at a cost of fundamentally changing the problem, making a comparison between methods impossible. Furthermore, a composition-only representation is inherently unsuitable for the direct prediction of most material properties that depend on the atomic structure. The structure-informed descriptors can include much more information related to interatomic interactions, making them more robust and more physics-relevant. They also, implicitly or explicitly, include symmetries present in the material, which can be used to predict certain properties, such as zero piezoelectric response, with high confidence. Furthermore, such descriptors often include extensive composition-based arguments within them [31], making it possible to both recognize patterns in the property coming from different chemical species occupying the same structure and structural effects in the case of a single composition.

At the same time, it is important to consider that physically existing materials are rarely described by a single atomic configuration, usually requiring considerations for defects and coexisting configurations. Thus, like a traditional DFT-based modeling, in order to reproduce real material behavior, a structure-informed model will often require utilization of a method such as CALPHAD [35,36]. One of such methods, recently developed by authors and named “zentropy theory” shows the potential to connect individual configurations to predict

macroscopic properties, such as colossal positive and negative thermal expansions [37].

In some cases however, investigating all configurations can be a very challenging task (e.g., for high entropy alloys), necessitating the use of an elemental-only model trained to give predictions assuming future observations to be consistent with the past ones [38].

The structure-informed representation which was the ground for the present work has been developed by Ward et al. based on information from the Voronoi tessellation of a crystal structure [31]. Ward’s descriptor contains 271 features that combine information from elemental properties of atoms, such as shell occurrences, with information about their local environments, such as coordination number or bond lengths to neighbors. This approach was demonstrated to work excellently when comprehensively compared to two previous approaches based on the Coulomb matrix (CM) [33] and on the partial radial distribution function (PRDF) [32], when trained on the same data from the Open Quantum Materials Database (OQMD) and with the same machine learning algorithm. A more detailed overview is given in 2.1.

Ward et al. used an automated Random Forest ML algorithm [31] set to a fully automatic parameter selection. While fairly common, that approach without complexity limit for the model, and when trained on over 400,000 materials, resulted in a forest composed of 100 trees with approximately 700,000 nodes each. Such model requires over 27 GB of RAM memory to run, making it unusable on a typical personal or lab computer. Such size also results in a relatively low efficiency, requiring over 100 ms to run on a high-performance lab computer [31].

In the present work, aforementioned issues are resolved through a targeted design of the ML algorithm to fully utilize the data and its representation. This is done by consideration of the problem formulation and the deep neural network technique (see 2.2), combined with iterative model design (see 2.5), and by designing and testing over 50 neural networks belonging to around 30 designed architectures. Notably, in the time between Ward’s work and the present paper, neural networks have been used in this application, e.g., [39], which uses residual neural networks. However, as we show in Section 3.2, the present paper provides more accurate predictions than both Ward’s model and the state-of-the-art neural network model [39].

Additionally, the present work brings two further improvements. The first one is good transfer learning ability, described in 3.5 allowing other researchers, at a relatively small cost, to adjust the model to small problem-specific databases, typically consisting of tens of DFT calculations or less. This method substantially improves predictions for similar materials while retaining the general knowledge learnt from the large dataset and demonstrates that the model learns features related to underlying physics. The second improvement is the end-user usability. While most of the materials-related ML model are reported in a reproducible way with an evaluation of the performance [31–33,40], only a fraction goes beyond to make models accessible to the community. The present work has been focused on creating a Findable, Accessible, Interoperable, and Reusable tool, inspired by FAIR principles [41], created open-source with common and convertible data formats as is described in more detail in 2.3. This lead to many standalone components combined into an end-user tool, described in 3.7, that is ready to use without any costly computation to create the model and can be run on any modern computer, as low-power as smartphones.

2. Methodology

2.1. Descriptor used

A descriptor of a material is a point in a well-defined multidimensional property space that can be used to represent knowledge associated with entries in a database in vector form. Within the present work, the property space has 271 dimensions (corresponding to 271 features) related to elemental properties and atomic structure of an arbitrary crystalline material, as designed by Ward et al. [30,31] utilizing the voro++ code [42]. These features can be categorized as:

Table 1
List of features with descriptions.

Site statistics	Difference statistics	Name	Description
1-4	–	Effective coordination number	Mean, mean abs error, min, max
5-7	–	Mean bond length	Mean abs error, min, max
8-11	–	Bond length variation	Mean, mean abs error, min, max
12	–	Cell volume variation	Variation in the voronoi cell volume no statistics
13-15	–	Mean WC magnitude	shells 1-3, global non-backtracking
16	–	Packing efficiency	no statistics
133-138	17-21	Atomic number	
139-144	22-26	Mendeleev number	
145-150	27-31	Atomic weight	
151-156	32-36	Melting temperature	
157-162	37-41	Column	Group in periodic table
163-168	42-46	Row	Period in periodic table
169-174	47-51	Covalent radius	
175-180	52-56	Electronegativity	
181-210	57-81	Valence electron count	Listed for s, p, d, f orbitals and total
211-240	82-106	Unfilled count	Number of unfilled orbitals Listed for s, p, d, f orbitals and total
241-246	107-111	Ground state volume	
247-252	112-116	Ground state band gap	
253-258	117-121	Ground state magnetic moment	
259-264	122-126	Space group number	Index of space group
127	–	Number of components	No statistics
128-132	–	ℓ^p -norms of component fractions	$p \in \{2, 3, 5, 7, 10\}$
265-268	–	Fraction of valence electrons in s, p, d, f orbitals	No statistics
269	–	Can form ionic compound	Boolean, no statistics
270-271	–	Ionic character	Max, mean over pairs of species

Site Statistics refers the mean, range, mean absolute error, maximum, minimum, and mode unless otherwise stated in the description. Difference Statistics refers to the mean, mean absolute error, minimum, maximum and range of the differences between neighboring sites in a structure, weighted by the size of the face between them in the Voronoi tessellation.

- **Elemental Attributes** (145 total): Attributes which only depend upon the elements present and their stoichiometry.
 - **Stoichiometric Attributes** (6): Describe the components fractions.
 - **Elemental Properties Attributes** (132): Contain statistics taken over the various elemental properties, weighted by the stoichiometry of the structure.
 - **Attributes based on Valence Orbital Occupation** (4): Depend upon the distribution of valence electrons across different orbitals, i.e. on the total number of valence electrons in each orbital across the structure.
 - **Ionic Character Attributes** (3): Attributes which encode whether the material is ionically bonded.
- **Structural Attributes** (126 total): Attributes which depend on the precise structural configuration, i.e. exactly how the atoms are arranged in space.
 - **Geometry Attributes** (16): Attributes which depend upon the spatial configuration of atoms only.
 - **Physical Property Differences Attributes** (110): Contain statistics taken over the differences between elemental properties of neighboring sites in the structure, weighted by the size of the Voronoi cell face between the neighbors.

A complete table list of features is given in [Table 1](#). Further details can be found in [\[30,31\]](#).

2.2. Machine learning techniques overview

This section gives a brief overview of the employed machine learning techniques and terminology, described in more detail in [Appendix A](#). The interest is placed on the statistical problem of regression, whose goal is to learn a functional relationship $f : X \rightarrow Y$ which minimizes the risk (also known as loss or expected error) [\[43\]](#) given by

$$R(f) = \mathbb{E}_{x,y \sim p} l(y, f(x)). \quad (1)$$

Here X denotes a space of input features, Y denotes an output space, the expectation above is taken over an unknown distribution P on $X \times Y$ (representing the true relationship between inputs and outputs), and l is a given loss function.

In the specific application considered here, the function f which is to be learned, maps input material structures (arrangements of atoms) x to the predicted formation energy y . The distribution P is unknown, but samples (x_i, y_i) are given, consisting of structures x_i and corresponding predictions y_i which are used to learning f . In the present case, this data comes from the OQMD and other smaller materials databases.

In order to learn the relationship f from the data, the empirical risk

$$L(f) = \frac{1}{n} \sum_{i=1}^n l(y_i, f(x_i)), \quad (2)$$

is minimized over a class of functions defined by a neural network architecture. A neural network architecture consists of a sequence of alternating linear functions and point-wise non-linear functions defined by an activation function (see [\[44\]](#) for more information about neural networks). As the loss function l in [\(2\)](#) the ℓ^1 -loss function $l(y, x) = |x - y|$ is used. The neural networks are trained on this loss [\(2\)](#) using the common ADAM optimizer [\[45\]](#).

An important issue when training complex statistical models is the overfitting, which occurs when a model accurately fits the training data but fails to generalize well to new examples. In order to detect overfitting, the standard practice of dividing the data into training, validation, and test datasets [\[46\]](#) is used. In order to mitigate overfitting, dropout [\[47\]](#) and weight decay, two standard methods for regularizing neural networks, are used. In [Section 3.1](#), [Fig. 3](#) illustrates overfitting mitigation effects on the training process of neural networks designed in the present paper.

2.3. Software used

The choice of software for the machine learning portion of this project was Apache MXNet [\[48\]](#) due to its open source nature, model

portability, and state-of-the-art scalability, allowing the same code to run on a laptop with a low-power CPU/GPU and a supercomputer (e.g., ORNL Summit) with hundreds of powerful GPU's. Its portability allows trained networks to be converted and used with other popular frameworks such as Google Tensorflow, PyTorch, or even Apple Core ML, making results of the present paper highly accessible.

MXNet framework was used through Wolfram and Python languages. Wolfram Language was used primarily for the network architecture design, training, and testing, as it provides an excellent interface with detailed training results shown in real-time during the training process. It also provides good out-of-the-box performance due to its well-optimized memory handling when training on a single GPU setup.

Python, on the other hand, was used when writing the end-user tool for running previously trained networks. This choice was made so that the software is completely open-source and can be easily reused for specific purposes or incorporated within other packages. Furthermore, Python allowed quick implementation of a Graphical User Interface (GUI) through the wxpython package.

2.4. Data acquisition and curation

Four sets of data were used within the present work. The largest by volume and significance was the Open Quantum Materials Database (OQMD) [1,2], which contains the results of DFT calculations performed by the Vienna Ab Initio Simulation Package (VASP) [49] for a broad spectrum of materials. The snapshot used here was extracted from the database by Ward et al. in 2017 and contained 435,792 unique compounds [31]. The choice of 2017 snapshot rather than the current one was made to ensure direct performance comparison between new and previously reported methods. The second database was a part of the Inorganic Crystal Structure Database (ICSD), a subset of the OQMD with only experimentally obtained structures containing around 30,000 entries. ICSD was primarily used for the quick design of simple neural network architectures at the beginning, and OQMD used for more complex models designed later.

Two smaller datasets were used, in addition to these large databases. The first small dataset contained DFT-calculated formation energies of Fe–Cr–Ni ternary σ -phase endmembers in the 5-sublattice model [50]. As this model contains 5 chemically distinct positions (Wyckoff positions), populated by one of 3 elements, in total it included 243 (3^5) structures with 30-atom basis each. This data served as an example of a relatively complex structure that was not included in the OQMD. Furthermore, it was a test case of a material that is highly industry-relevant, as it causes steel embrittlement [51] and is costly to investigate using traditional methods due to compositional and configurational complexity. The second small dataset included 13 Special Quasirandom Structures (SQS), which are the best periodic supercell approximations to the true disordered state of metal alloys [52–54]. SQS structures in this set were binary alloys containing Fe, Ni, Co, and V, laying on deformed FCC (A1), BCC (A2), or HCP (A3) lattices. The main purpose of these smaller datasets was to test the performance in extrapolation from OQMD, in a particular case of interest for the author's.

During the network design process described in 2.5, it was found that a small fraction of the OQMD dataset (under 0.03%) contains anomalous values of formation energy above 10 eV/atom. In the extreme case of CuO_2 (OQMD ID: 647358) this value was 1123 eV/atom or 108350 kJ/mole. Since the source database contains hundreds of thousands of data points reported by many scientists, it can be expected that a small fraction of the data may contain some sort of errors and typos. In the present work, they were removed from all datasets used for training and evaluation.

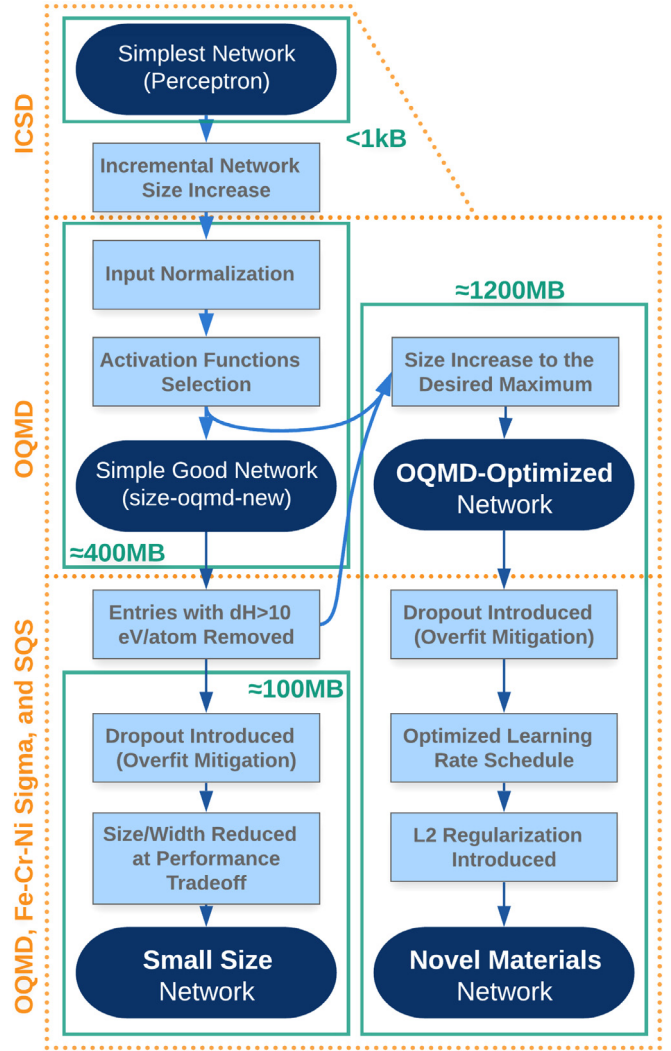


Fig. 1. The model design process schematic.

2.5. Neural network design process

This section conceptually outlines the network design process leading to the final models. All essential details regarding the design and performance of intermediate models, useful for better understanding changes and for applying the similar approach in different problems, can be found in Appendix B.

The design started with the simplest single-layer neural network (perceptron) with the Sigmoid activation function, trained on the ICSD and its smaller subset, to provide a baseline for the design. Then, the process was conducted in the following steps:

1. The network size has been increased step-wise while training on the ICSD dataset (30k+ entries). Results were extrapolated to estimate network size suitable for larger OQMD (400k+) to be 4 hidden layers in a (10000, 10000, 1000, 100) configuration.
2. To improve convergence during the training, descriptor features values were normalized to their maximum values present in the OQMD dataset.
3. Performance and time to convergence were improved by moving from Sigmoid activation function to a mix of Soft Sign, Exponential Linear Unit, and Sigmoid. This relatively simple model has improved performance over the existing Random Forest model [31], achieving MAE of 42 meV/atom on the same dataset.

4. At this step, it was noticed that a small fraction (around 0.03%) of data points exhibits extreme errors, as high as over 1,000,000 meV/atom causing some instability during the training process, despite the large batch size of 2048. They also caused a high deviation in test MAE values across repeated model training rounds. As describes in 2.4, these were identified to be a few rare errors in the dataset and removed during later model design.

5. The network size was increased to around 1 GB limit (maximum size target) by the addition of two more 10,000-width layers. This **OQMD-optimized** network has achieved the best performance on the OQMD out of all designed in the present paper, with an MAE of 28 meV/atom. Performance analysis can be found in 3.2 and in Fig. 5.

6. After the good performance on the OQMD was achieved, the design goals shifted to (1) reducing the training-set-to-validation-set error mismatch during the network training, while (2) keeping the test MAE on the OQMD on a suitable level (below 50 meV/atom), and (3) improving performance on datasets not presented to network before (see 2.4). The first step was the introduction of Dropout layers [47], described in more detail in Appendix A, which allow for better distribution of knowledge across the network.

7. The introduction of strong Dropout [47] made the network prone to falling in local minima, which was solved by the introduction of a changing learning rate schedule.

8. With optimized network architecture, lastly, the descriptor interpretation by the network has been modified through the introduction of L2 regularization [55], a technique which assigns an error penalty for “attention” (input layer weights) to each of the descriptor features, effectively refining features in the descriptor to only the most significant ones. Fig. 4 ranks them. The resulting **Novel Materials** model achieved a much lower training-set-to-validation-set error mismatch (1.15 vs 1.57 after 240 rounds), presented in Fig. 3 as a function of training progress. On the OQMD test set, it achieved a higher, yet suitable 49 meV/atom.

9. To cater to applications requiring very high throughput or low memory consumption, an additional **Small Size** network was designed by adding Dropout to one of the earlier networks, designed before the size increase step, and then reducing its size to the desired level. It was found that after reduction of total size from around 400 MB to around 100 MB, the network retained MAE of 42 meV/atom on an OQMD test set and further reduction was possible if needed for the application.

3. Results

3.1. Final predictive models

Throughout the architecture design process described in 2.5, detailed in Appendix B, and depicted in Fig. 1, new networks were designed and tested in various ways, leading to about 50 predictive models (trained neural networks) with varying training parameters and training data. The majority of the intermediate networks were stored for the record, and are available upon request. Details regarding hyper-parameters and training routines used to obtain three resulting models can be found in Appendix A.

Out of all trained neural networks, three were selected and can be considered final outcomes of the design process, optimized for different objectives. Their architectures are presented in Fig. 2. The first one, denoted NN9, was created specifically for the OQMD performance. This was the same objective as in the study by Ward et al. [31] and its performance serves as a direct comparison to the Random Forest method employed in that paper [31] and other works [32,33].




Fig. 3. Training Loss to Validation Loss in a model that does without (NN9) and with overfitting mitigation (NN20), plotted versus training progress.

Input	vector (size: 271)	Input	vector (size: 271)	Input	vector (size: 271)
1 LinearLayer	vector (size: 10000)	1 LinearLayer	vector (size: 10000)	1 LinearLayer	vector (size: 4000)
2 SoftSign[x]	vector (size: 10000)	2 SoftSign[x]	vector (size: 10000)	2 SoftSign[x]	vector (size: 4000)
3 LinearLayer	vector (size: 10000)	3 DropoutLayer	vector (size: 10000)	3 DropoutLayer	vector (size: 4000)
4 LogisticSigmoid	vector (size: 10000)	4 LinearLayer	vector (size: 10000)	4 LinearLayer	vector (size: 4000)
5 LinearLayer	vector (size: 10000)	5 LogisticSigmoid	vector (size: 10000)	5 LogisticSigmoid	vector (size: 4000)
6 LogisticSigmoid	vector (size: 10000)	6 DropoutLayer	vector (size: 10000)	6 DropoutLayer	vector (size: 4000)
7 LinearLayer	vector (size: 10000)	7 LinearLayer	vector (size: 10000)	7 LinearLayer	vector (size: 4000)
8 LogisticSigmoid	vector (size: 10000)	8 ExponentialLinearUnit[x]	vector (size: 10000)	8 ExponentialLinearUnit[x]	vector (size: 4000)
9 LinearLayer	vector (size: 1000)	9 DropoutLayer	vector (size: 10000)	9 DropoutLayer	vector (size: 4000)
10 ExponentialLinearUnit[x]	vector (size: 1000)	10 LinearLayer	vector (size: 10000)	10 LinearLayer	vector (size: 1000)
11 LinearLayer	vector (size: 100)	11 ExponentialLinearUnit[x]	vector (size: 1000)	11 ExponentialLinearUnit[x]	vector (size: 1000)
12 ExponentialLinearUnit[x]	vector (size: 100)	12 DropoutLayer	vector (size: 1000)	12 DropoutLayer	vector (size: 1000)
13 LinearLayer	vector (size: 1)	13 LinearLayer	vector (size: 100)	13 LinearLayer	vector (size: 100)
Output	vector (size: 1)	14 ExponentialLinearUnit[x]	vector (size: 100)	14 ExponentialLinearUnit[x]	vector (size: 100)
		15 LinearLayer	vector (size: 1)	15 LinearLayer	vector (size: 1)
		Output	vector (size: 1)	Output	vector (size: 1)

Fig. 2. Three selected architectures designed within the present work. Optimized for: (Left) OQMD performance, (Middle) predicting new materials, (Right) small size at good performance. Internally in the code, they are designated as NN9, NN20, and NN24.




Fig. 4. Distribution of sums of squared input weights. High values correspond to attributes that were not lowered due to their contribution to pattern recognition of the model. 15 attributes with the highest values are labeled. The labels are taken from the descriptor definition in [30].

The second network was optimized for improved pattern recognition on OQMD and improved performance on non-OQMD datasets used in the present work (i.e. SQS/ σ -phase datasets). This was achieved primarily through extensive overfitting mitigation, applied during design and training (see Fig. 3), which leads to a network with improved generalization/materials-discovery capability. Furthermore, one of the overfitting mitigation methods, namely the regularization described in 2.2, have allowed identification of descriptor attributes that contributed the most to the predictive capability and the ones that were almost completely discarded once the penalty for considering them was assigned. Fig. 4 presents the distribution of sums of squared weights between each neuron in the input layer (each of the 273 descriptor features) and all 10,000 neurons in the first hidden layer.

Feature rankings, such as presented in Fig. 4, allow a more efficient selection of input features in future studies looking into the same problem; thus both reducing the number of features that need to be computed for each atomic configuration and the total number of weights in the network. Furthermore, it can be used to gain an

insight into the model interpretability. Looking at the specific ranking for NN20, the high-impact features present a mix of elemental features, likely allowing the model to establish some formation energy baseline for a given composition, and structure-informed features allowing to distinguish between polymorphic configurations. High impact elemental features include different statistics on elemental melting temperatures and ground-state structure volume per atom. The structural features extend them by considering how they differ between neighboring atoms and also include purely structural features such as packing efficiency and variance in Wigner–Seitz cells volumes. A complete ranking of features is included in Appendix C.

The third network, denoted NN24, was created for memory/power-constrained applications requiring a balance between OQMD performance and memory intensity and processing power required. Model parameters contained in this architecture occupy only 145 MB, over 8 times less than two other models and around 200 times less than the model reported by Ward et al. [31].

3.2. OQMD data performance

As described in 2.5, all three final networks were evaluated on a randomly selected subset of the OQMD to give a comparison between the state-of-the-art model presented by Ward et al. [31] and the present ML method. This random subset consisted of 21,800 OQMD entries, constituting approximately 5%, which were not presented to the network, nor used for evaluation at any stage of the training process. This sample size was considered to be representative of the whole dataset once the small fraction (0.026%) of likely incorrect entries were removed from the dataset as described in 2.4. The random selection itself was initially performed separately for each training process and recorded after completion. Later, when networks were modified to mitigate overfitting, a single random subset was used for all of them to allow more careful design and more accurate comparative analysis of results. Fig. 5 gives (1) prediction vs OQMD values of formation energy plot, (2) statistics related to the error in predictions relative to the OQMD values, and (3) a histogram of the absolute error in predictions relative to the OQMD values.




Fig. 5. Performance of 3 selected neural networks on a random subset of 21,800 entries from OQMD. (Left) OQMD performance, (Middle) predicting new materials, (Right) small size at good performance. Internally in the code, they are designated as NN9, NN20, and NN24.

Table 2

Comparison of our method with existing state-of-the-art methods.

Method	Formation energy MAE	Convex hull MAE
SIPFENN (This work)	28.0 meV/atom (OQMD Opt.)	32 meV/atom (Novel. Mat.)
Ward2017 [30,31]	80 meV/at	N/M
ElemNet [29]	N/A	50 meV/at
IRNet [39]	38 meV	N/M
Roost [34]	N/A	29 meV/at – 24 meV/at

N/A and N/M respectively stand for not applicable and not measured.

3.3. Existing methods comparison

In this section, the performance of the models is compared with a few similar existing approaches based on the OQMD dataset, when formation energy of a structure is predicted [30,31,39], or its subset of the convex-hull structures, when formation energy of the most stable structure is predicted [29,34]. This division is made based on the reasoning presented in 1. While the latter type cannot be used to predict the formation energy of any arbitrary structure, the structure-informed models like SIPFENN (the present work) can be tested on the convex hull structures.

The results are shown in Table 2. The SIPFENN convex hull MAE has been reported based on using the Novel Materials Model limiting the original test set to structures laying within 50 meV/atom from the convex hull. From these results, we can see that the SIPFENN neural networks approach outperforms existing state-of-the-art methods for predicting the formation energy of any material. At the same time, while not being the best, it is capable of reaching performance levels of specialized models in predicting the formation energies of structures laying on the convex hull.

3.4. Non-OQMD data performance

Models created in the present work, specifically the ones optimized for predicting the formation energy of new materials, were designed and implemented to serve as tools for materials discovery. Evaluating their performance on data from the same source as the training set done in 3.2 is inherently biased towards favoring models that provide the best fit to the prior (training) knowledge. This is amplified by the fact that many entries in the database are reported in groups that come from common studies and span similar materials, causing high domain clustering, which in some cases effectively makes such evaluation more akin to interpolation than extrapolation of knowledge.

To partially mitigate the described issue, the performance of the models was also evaluated on two smaller non-OQMD databases, described in 2.4, representing an example of chemistries and structures

that were of interest to the authors project on Ni-based superalloys. At the same time, they were not directly presented to the network in any capacity during the training process.

In all cases, models created in the present paper were able to achieve approximately the same performance as on a random selection from the OQMD. To give a more in-depth analysis of the results, Fig. 6 shows a magnified view of the predictions and basic statistics on the agreement between predictions and the database for the three models developed in the present work.

While all three models performed at around the same MAE level as for the OQMD, network optimized for new materials, the NN20 and NN24, performed better in the non-OQMD test cases of interest, providing major increases in correlations, significant for ranking of end-member configurations, except for 4 SQS configurations which were underestimated. The Pearson correlation slightly decreased in the first case and slightly increased in the second case. In both cases, the mean absolute error decreased by about 20% compared to the OQMD-optimized model.

3.5. Transfer learning capability

In this section, the technique of transfer learning is considered. It has been observed among deep learning models across a variety of domains [56–59] and refers to the ability of properly trained deep learning models to ‘transfer’ their knowledge to related tasks. In the least complex approach, one does this by simply ‘fine-tuning’ the parameters of the model using new training data (from the new task). This methodology has shown in practice that deep neural networks are often able to transfer knowledge between different but related tasks. Such a problem is analogous to many others in materials science, where general knowledge is used to make meaningful statements without statistically significant patterns in locally available data.

It is shown that a network trained on the OQMD database, which covers a broad yet limited spectrum of materials, can be quickly adjusted to materials outside of this spectrum with very little additional cost relative to the initial training. Specifically, the transfer learning capability of a network trained in this way on the set of all (243) Fe–Ni–Cr σ -phase 5-sublattice model endmembers, described in 2.4, was tested. The ML model was first trained on a broad and general material dataset (OQMD) and then further trained (i.e., re-trained) for a given number of rounds on the new data (Fe–Ni–Cr σ -phase dataset) to adapt to the new system, while still conserving its broad knowledge, and can be thought of as fine-tuning a model to improve extrapolation outside of a prior knowledge space.

In order to achieve good performance, both the number of rounds and the learning rate have to be optimized. This can be accomplished by investigating the dependence of error on the fraction of available




Fig. 6. Performance of 3 selected neural networks on non-OQMD data described in 2.4. Evaluated on (red) Fe–Cr–Ni σ -phase and (blue) SQS dataset. Networks organized by columns; optimized for (left) OQMD performance, (middle) predicting new materials, (right) size-constrained applications. Internally in the code, they are designated as NN9, NN20, and NN24 respectively.




Fig. 7. MAE evolution of NN20 model re-trained for 25 additional rounds on an increasing fraction of data from Fe–Cr–Ni σ -dataset. Presents the dependence of transfer learning from new data for different learning rates expressed as fractions of default ADAM learning rate (0.001).

data while one of these parameters is fixed. Fig. 7 presents the dependence of transfer learning from new data for different learning rates expressed as fractions of default ADAM learning rate (0.001 shared across a vast majority of software).

As shown, in this case, the default learning rate (100%) cannot be used for the transfer learning as it will adjust network parameters in both an unreliable and detrimental fashion, resulting in poor

performance on the whole system of interest (both training and test sets) as shown in Fig. 7. The same behavior would be observed if the process were conducted using an automated model design available in software such as MATLAB or Mathematica. The 10% learning rate provided reliable enough outcomes and allowed a better performance improvement given little data, relative to using a 1% learning rate (relative to the default). The second parameter to be optimized was the number of re-training rounds, as presented in Fig. 8.

Fig. 8 shows that use of too few retraining rounds causes unreliable outcomes, while too many causes overfitting for low amounts of new data. In the case of Fe–Cr–Ni σ -dataset, retraining for 10 or 25 rounds provides balanced results across the whole dataset. With parameters for the process set to 10% learning rate and 25 additional rounds, the performance can be evaluated graphically, as presented in Fig. 9.

As depicted, adding just 10% of DFT-calculated data (24/243 end-members) provided a significant improvement in the prediction quality over the system, including the other 90% was never shown to the model. This result indicates that the models in the present paper can be combined with partial data obtained through DFT calculations to create accurate predictive tools for a specific closed material system, such as sublattice endmembers, and potentially limit the number of calculations required within the study. This can then provide the ability to investigate broader material search spaces at a given computational cost.

Furthermore, the presented transfer learning capability could be used for a more broad materials exploration without a well-defined finite search space like the ternary Fe–Cr–Ni σ -phase. In such a case,

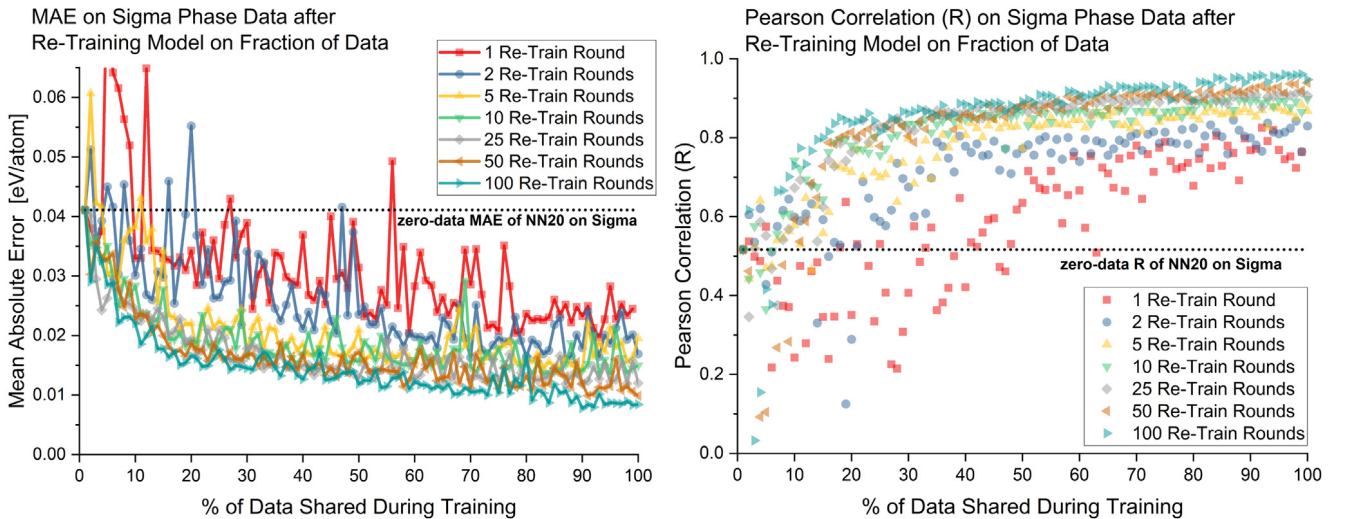


Fig. 8. MAE and Pearson correlation (R) evolution of NN20 model re-trained at 10% learning rate on an increasing fraction of data from Fe–Cr–Ni σ -dataset. Presents the dependence of transfer learning from new data for different re-training rounds numbers.

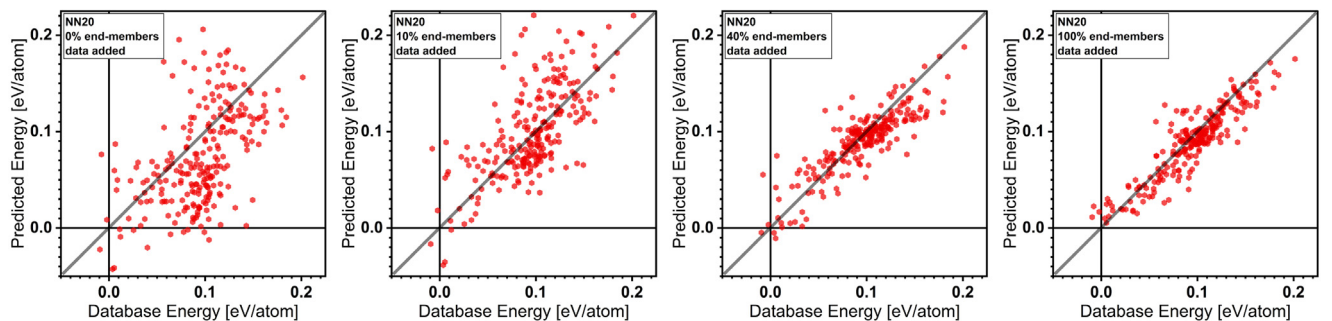


Fig. 9. Performance of a new-materials-optimized network (NN20) on σ -phase data. Left-to-right: as trained on the OQMD, with additional training on 10%, 40%, and 100% of the Fe–Cr–Ni σ -phase end-member data. The points on the figure correspond to all end-members (both training and testing data). Corresponding MAE and R are presented in Fig. 8 (gray rhombus points).




Fig. 10. MAE of predictions evaluated on test set data vs number of newly available training datapoints. 1180 blue points correspond to single transfer learning processes. Red plot gives mean MAE and standard deviation. Both plots contain the same data.

it is better to evaluate and report the performance of the model on a test set that was not presented during the training and report, as a function of the number of added data points (new DFT calculations). With such a problem statement, the transfer learning process has been repeated 1180 for the statistical significance of the outcomes, which are presented in Fig. 10.

As presented in Fig. 10, adding just a small number of new data points allows to nearly half the MAE (around 20 datapoints). Furthermore, evident from the right plot, the mean performance increase is on average linear in log-lin scale and highly predictable ($R^2 = 0.98$).

3.6. Model limitations

As with any modeling tool, this modeling effort has some inherent limitations, coming from both data and methods used to create it. The most significant one comes from the type of data used for training of the model, where all data points correspond to DFT-relaxed structures, sitting in local minima in the configuration energy landscape. Thus, all energy predictions are given under an assumption that the input structure is fully relaxed with DFT settings inherited from the OQMD database [1]. At the same time, since the model was trained on many local energy minima configurations analyzed on the level of single-atom chemical environments, it should be able to approximate values for unrelaxed structures based on substitution from prototypes or similar compounds. Testing of this is performed by Ward 2017 [31], where it is shown that (a) in most of the test cases, the before-after relaxation energy difference is negligible in comparison to the DFT-ML difference for Ward 2017 model and usually much lower than the test MAE for models discussed in this work, and (b) in some particular cases ($\text{Li}_6\text{CaCeO}_6$) can be very high.

When faced with a new configuration, the model can thus either be used to (1) give an accurate prediction if the configuration is already relaxed or (2) give an approximate result that needs to be validated with DFT if confidence in the result is needed. This is inherent to all structure-informed ML models. One possible solution to partially mitigate this limitation is to perform relaxation using the model, which should work reasonably well for most materials.

Discussion of such relaxation procedure in detail is extensive and beyond the scope of this work, yet a preliminary approach was constructed using the Novel Material Model (NN2) and deployed on all 16 end-members of Pd-Zn γ -brass crystal structure in an iterative fashion. At each iteration, first, the local energy gradient for each atom was calculated by comparing the starting configuration with perturbations in x, y, z directions. Then, all atoms were displaced proportionally to the gradient in 100 discrete steps, reaching some local minimum, which acted as a starting point for the next iteration. An example for Pd_8Zn_5 is presented in Fig. 11.

As shown in 11, the resulting relaxation reduced predicted formation energy by 4 meV/atom for this particular end-member. In




Fig. 11. Local energy landscape relaxation of Pd_8Zn_5 γ -brass crystal structure guided by Novel Material Model (NN20).

the other 15 cases, results were similar, ranging between near 0 and 15 meV/atom, converging into fine local minima, expected to correspond with true local relaxations; however, extensive research into the problem is needed before conclusions can be drawn.

3.7. End-user implementation - SIPFENN

One of the main objectives of the present paper was to create a tool that is transparent, easy to use by the research community, and easily modifiable. This led to the creation of SIPFENN (Structure-Informed Prediction of Formation Energy using Neural Networks) software. SIPFENN provides the user with near-instant access to the models presented in 3.1. In the future, this selection will likely be further expanded. On the user side, the use of the software is as easy as selecting one of the models, specifying a folder containing structure information files like POSCARs [60] or CIFs [61], running the predictions, and saving results.

SIPFENN was written entirely in Python to allow other researchers to easily modify it and adjust it to specific needs. Its schematic of operation is presented in Fig. 12. In broad scope, it first performs the structure analysis and modifications using the Python Materials Genomics library (pymatgen) [62]. In the current implementation, it imports all structure files, analyzes the stoichiometry, creates unique names based on that, and exports them as POSCAR files. This is a rather simple task, however pymatgen is a powerful tool with a suite of more complex analytical tools that can be quickly implemented into SIPFENN by the user with even basic Python skills. Following

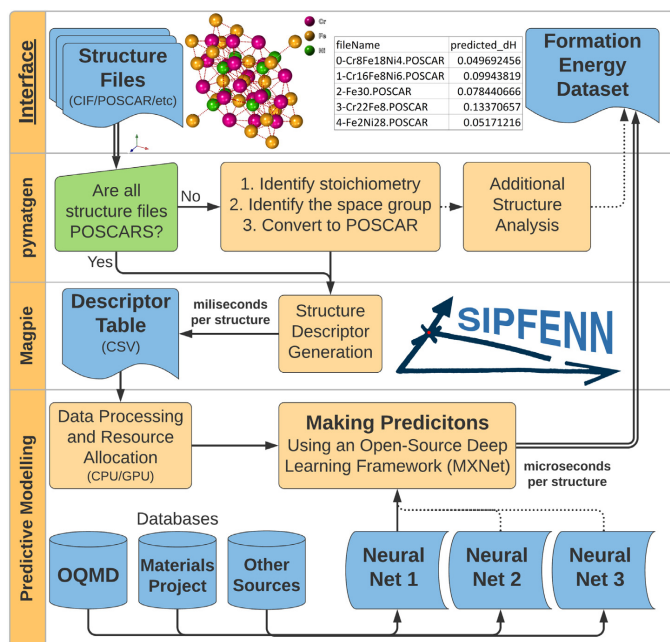


Fig. 12. SIPFENN schematic description of operation.

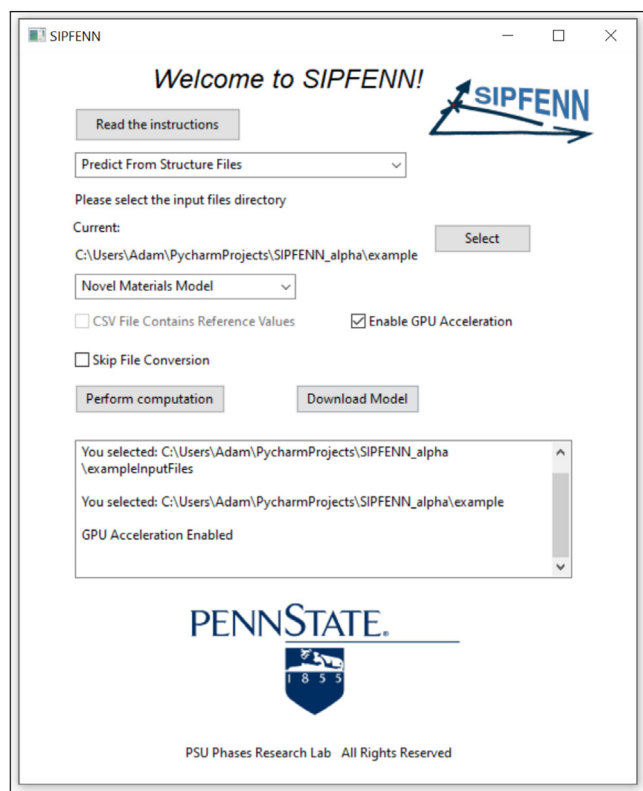


Fig. 13. A snapshot of the graphical user interface of SIPFENN.

the analysis, SIPFENN runs java-based Magpie [30] which calculates a descriptor for every imported structure and exports the result as a CSV file. This file is a descriptor table, where each row corresponds to a single material, and which can be stored and re-used later to run multiple predictive models at a fraction of the original computation time. It can also be used to create datasets for training procedures by

replacing the last column with calculated or experimental values of formation energy.

Finally, the descriptor table is imported into the MXNet library framework, allocated into the CPU or GPU memory based on user selection, and evaluated using the selected predictive model. Once results are obtained, they are exported in CSV format and can be analyzed by any spreadsheet software such as Microsoft Excel.

SIPFENN was planned as a command-line tool, however, it was recognized that some users, especially those with little computational background, may find that difficult. Therefore, a simple graphical user interface (GUI) was created using wxPython library. It incorporates all the capabilities of the command line version. Furthermore, it lets the user download the predictive models from a repository in a single click. A sample snapshot of the GUI before performing calculations is presented in Fig. 13.

4. Conclusions

In the present paper new machine learning models and a ready-to-use tool were created, based on the dataset and descriptor design by Ward et al. [31]. Models reported in this work significantly improve upon existing methods, both in terms of performance and accessibility. For the most direct comparison, one of the designed models has been optimized for performing well on a random subset of the OQMD database and achieved an MAE of 28 meV/atom, compared to 80 meV/atom in the original Ward et al. paper [31], and to 38 meV/atom in the most recent model called IRNet [39]. Furthermore, it was shown that the error of the model is lowered when applied to the problem of finding the convex hull energy, achieving levels comparable with the current state-of-the-art approaches [29,34].

In addition, using appropriate overfitting mitigation efforts, such as Dropout and L2 regularization, models tuned for generalization to other types of materials datasets were developed. To test this, the models were evaluated on two datasets not contained within the OQMD, namely all end-members (243) of 5-sublattice topologically-close-packed Fe–Cr–Ni Sigma-phase [50,51] and a few selected random-solution-approximating SQS [52–54]. The MAE values for these two test sets were found to be close to the values obtained on a test set from the OQMD. This exemplifies that the models are able to generalize to new datasets.

Furthermore, it was shown that models created within the present paper can be used for transfer learning, where vast knowledge of a broad spectrum of materials is combined with as little as a few DFT-datapoints from a specific materials system to provide excellent results within that specific system. Such at least partially process mitigates the issue of low data availability, present in numerous materials science problems, and consequently allows users to investigate a broader scope of materials at the same computational cost.

Finally, the three neural network models designed within the present paper were used, in conjunction with additional software, to create an end-user tool called SIPFENN. SIPFENN's capabilities extend far beyond allowing validation of the presented results. It is implemented to work without any intensive computations on the user side, using models accessible from a repository, requiring only a quick one-click model download to run. It is very fast thanks to using one of the industry's leading ML frameworks capable of well-optimized computations on GPUs. Furthermore, it is an open-source tool written in Python, which can be easily modified to specific needs in a straightforward way without extensive changes in the code.

CRedit authorship contribution statement

Adam M. Krajewski: Conceptualization, Methodology, Software, Validation, Visualization, Data curation, Writing – original draft. **Jonathan W. Siegel:** Conceptualization, Methodology, Software, Validation, Writing – original draft. **Jinchao Xu:** Funding acquisition, Writing – review & editing. **Zi-Kui Liu:** Supervision, Funding acquisition, Project administration, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The present work was financially supported by the ICDS Seed Grant from the Pennsylvania State University, the Office of Naval Research (ONR) via Contract No. N00014-21-1-2608, the National Science Foundation (NSF) via Grants No. CMMI-1825538 and CMMI-2050069, and the Department of Energy (DOE) via Award Nos. DE-FE0031553, DE-EE0008456, DE-NE0008945, and DE-AR0001435.

We would like to thank Zhengqi Liu for his help implementing the graphical user interface, Dr. ShunLi Shang for providing the Fe–Cr–Ni σ -phase dataset, and Brandon Bocklund for providing the SQS dataset.

Software and data availability

The most recent version of SIFPENN code is available through Penn State's Phases Research Lab website at www.phaseslab.com/sipfenn in (1) a minimal version that can be run on pre-computed descriptors in CSV format as well as (2) ready-to-use version with pre-compiled Magpie [30]. SIFPENN contains hard-coded links to neural networks stored in the cloud that can be downloaded at a single-click (see Fig. 13).

All neural networks are stored in both (1) open-source MXNet format maintained by Apache Foundation and used within SIFPENN, and in (2) closed-source WNet format maintained by Wolfram Research and having the advantage of even easier deployment, as well as guaranteed forward compatibility with future versions of Wolfram Language.

For ensured longevity of results, SIFPENN neural networks are stored through the courtesy of Zenodo.org service under <https://doi.org/10.5281/zenodo.6133037> at the CERN's Data Centre.

Appendix A. Machine learning overview

The class of deep learning methods has been remarkably successful in recent years in applications ranging from computer vision to natural language processing and simulations of quantum systems [63–66]. Although deep neural networks have existed for a long time [67], and had been successfully applied to computer vision tasks [68–70], a major breakthrough was the AlexNet network [71], which dramatically improved the accuracy achievable on large-scale image classification. Following this success, deep neural networks have been very intensively studied and applied to a variety of problems [63–65]. Deep neural networks are particularly effective when applied to regression problems, where one is learning a functional relationship between a feature and a prediction. For many problems, deep neural networks are able to achieve significantly better performance than competing machine learning methods, due to their ability to learn more complex relationships. With materials science being a field where many complex dependencies are intertwined, it is to be expected that this superior pattern recognition can carry over to the improvement in the prediction of material properties.

A.1. Regression problem formulation and artificial neural networks

The general formulation of a regression problem in statistical machine learning is to find a function $f : X \rightarrow Y$ which minimizes the risk [43], also known as loss or expected error.

$$R(f) = \mathbb{E}_{x,y \sim \mathcal{P}} l(y, f(x)). \quad (3)$$

Here X denotes a space of input features, Y denotes an output space, the expectation above is taken over an unknown distribution \mathcal{P} on $X \times Y$ (representing the true relationship between inputs and outputs), and l is a given loss function. The goal is to find a function f which accurately predicts the (potentially random) output y given an input x .

In the present work, $x \in X$ represents the input features (descriptor) characteristic of the material, and $y \in Y$ represents the formation energy. The distribution \mathcal{P} represents the true material–property relationship between given descriptor x and corresponding formation energy. This relation may not be as simple as mapping a given structure to an energy since different DFT methodologies may give different results, based on many variables, such as employed functionals. [72,73] Consequently it is useful to describe this relationship via a probability distribution. Furthermore, the loss function considered in the present paper is the commonly used ℓ^1 or absolute error (AE) loss function $l(y_1, y_2) = |y_1 - y_2|$.

In practice, the distribution \mathcal{P} is not known. Indeed it is this relationship that one is trying to learn in the first place. Instead, what is available is data $\{(y_i, x_i)\}_{i=1}^n$, which is sampled from \mathcal{P} . From this one forms the empirical risk [46,74]

$$L(f) = \frac{1}{n} \sum_{i=1}^n l(y_i, f(x_i)), \quad (4)$$

and seeks a function f which minimizes the empirical risk, also known as the training error.

In addition, one must specify the type of relationship that is expected to be found between the inputs $x_i \in X$ and the predictions $y_i \in Y$. This is done by restricting the function f to a specific class. For instance, by restricting f to be linear, which corresponds to looking for a linear relationship between x_i and y_i , one obtains a linear regression. On the other hand, choosing \mathcal{F} to be a reproducing kernel Hilbert space of functions on X with the same loss l one obtains the kernel ridge regression method. Thus in order to fit the model, the training error is minimized over a specific class of function \mathcal{F} , i.e. one solves the optimization problem

$$f^* = \arg \min_{f \in \mathcal{F}} L(f) = \arg \min_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n l(y_i, f(x_i)). \quad (5)$$

In this the class of functions \mathcal{F} is chosen as the set of functions defined by a neural network architecture (schematic in Fig. 14), which leads to a deep learning method. A neural network architecture consists of a sequence of alternating linear functions and point-wise non-linear functions [44]. In the Fig. 14 the nodes, or neurons, represent applications of a point-wise non-linear function, called an activation function, and the connections between nodes represent linear functions from the output of the nodes in one layer to the input of the next layer.

The class of functions represented by the neural network consists of the functions obtained by substituting different linear maps between each layer. Specifically, given weight matrices W_1, \dots, W_n and biases b_1, \dots, b_n , which are parameters of the network, the corresponding neural network function is given by the composition

$$f_{W_1, \dots, W_n, b_1, \dots, b_n}(x) = W_n \cdots \sigma(W_3 \sigma(W_2 \sigma(W_1 x + b_1) + b_2) + b_3) \cdots + b_n \quad (6)$$

where σ , called the activation function, is applied pointwise to each entry of the vector input (previous layer output). The neural network architecture is determined by the type, dimensionality, activation function σ , and arrangement of intermediate layers. This can potentially introduce some additional restrictions on the linear maps W_i , see for instance convolutional neural networks, where the linear maps W_i are restricted to be convolutions with small kernels [68,70,71].

Once the neural network architecture has been set, one must fit the values of the parameters W_1, \dots, W_n and b_1, \dots, b_n by optimizing the training loss L ,

$$\arg \min_{W_1, \dots, W_n, b_1, \dots, b_n} L(f_{W_1, \dots, W_n, b_1, \dots, b_n}). \quad (7)$$




Fig. 14. Simplified artificial neural network schematic.




Fig. 15. A schematic of overfitting. The overfit model (yellow) is too complex and memorizes the training data. This results in very low training error, but also very poor performance when predicting hidden data (test error) that follows the underlying phenomena (red).

This optimization problem is typically solved using stochastic gradient descent [70], or a more robust method such as ADAM [45], which was used in the present work. To solve the problem faster and to mitigate overfitting, which is discussed in the next sections, these methods form an estimate of the loss function gradient by considering a small subset of the data, called a batch. Each training step is done over all of the data in the batch, so parameters (w and b) are updated based on many data points, rather than a single one. Most of the models created in the present work used a batch size of 2048 data points.

This methodology has been successfully applied to a variety of practical machine learning problems [71,75,76]. Specifically relevant to the present work, neural networks have been applied to problems in computational materials science [77,78]. For example, in [77] neural networks are used to classify the phases of high-entropy alloys. For this application, their neural network models compare favorably to other machine learning algorithms such as k -nearest neighbor (KNN) and support vector machines (SVM). Furthermore, in [78] it is shown that even when training on small datasets which are typical of certain materials science problems, specifically in the prediction of solidification defects from optical microscopy data, deep neural networks can achieve better performance than other machine learning models. This is enabled by using a stacked auto-encoder (shallow neural network) to pre-train the deep neural network, whose weights are then fine-tuned on the small dataset. The present work complements these studies by applying deep neural networks to the prediction of thermodynamic quantities from atomic structure descriptors.

A.2. Overfitting and its mitigation

A major problem in statistical learning is avoiding overfitting [46], which, in simple terms, signifies that the model memorizes the training data instead of learning the true relationship between descriptors x and predictions y . This occurs when the class of functions \mathcal{F} is too

large, and at the optimal function f^* in (5) the empirical (4) and true risk (1) diverge sharply. This results in very low training error, but poor performance on data that was not presented to the network. (See Fig. 15.) Overfitting is typically detected by separating the training data into two sets, the data used in (5) to learn the function f^* , called the training data, and a separate set of data used to evaluate the performance of f^* , called the validation set. Consequently, in addition to the training loss in (5), the validation error

$$L_{val} = \frac{1}{m} \sum_{i=1}^m l(\tilde{y}_i, f(\tilde{x}_i)), \quad (8)$$

where $(\tilde{y}_i, \tilde{x}_i)$ for $i = 1, \dots, m$ is the validation set, which was not presented to the network when adjusting its parameters, is used to detect overfitting. The fraction of the dataset aside for validation set should be large enough to be representative of the whole dataset to provide statistically significant conclusions, yet small enough so that knowledge loss in the process is minimized. In the present work, a randomly selected 15% of every dataset has been used as validation sets for all training. This corresponded to 65,300 data points in the case of the OQDM dataset described in 2.4.

Typically, the validation loss will be greater than the training loss, as the validation set is not available for training. This is illustrated in Fig. 16, where the ratio between the validation loss (8) and test loss (5) during the course of two trainings of similar NN architectures on the same data with the same learning rate schedule has been plotted. This figure indicates that as the training proceeds, the gap between the training and validation errors widens and then increases. The size of this gap is an estimated measure of how much the model has overfitted to the data. In one of the models in this figure, extensive techniques to mitigate overfitting have been used, and for this model, the figure shows that the rate at which the model overfits to the data is much lower. At the same time both models exhibit similar performance on the test set.

There are numerous techniques used to prevent the issue of overfitting [46,79]. These include utilization of a regularization term $\lambda R(\theta)$ added to the training error (5) to give the regularized empirical loss function

$$f^* = \arg \min_{f \in \mathcal{F}} R_{emp}(f) + \lambda R(\theta). \quad (9)$$

A standard regularizer typically added to the linear regression is the ℓ^2 -norm $R(\theta) = \|\theta\|_2^2$, which is often called Tikhonov regularization [80] or ridge regression [81]. The ℓ^2 -norm is also a popular regularizer in deep learning problems, where it is referred to as weight decay [44]. In the context of the present work, it is implemented as a part of the training process, rather than network architecture, and causes rejection of some features in the descriptor that are not contributing to pattern recognition. Results of its implementation are shown throughout Section 2.5.




Fig. 16. Training loss to validation loss in a model that does without (NN9) and with overfitting mitigation (NN20), plotted versus training progress.

Another important method used to prevent overfitting in machine learning is the Dropout technique [47]. The concept behind Dropout is to prevent neurons in the network from becoming overly dependent on the output from a specific neuron in the previous layer, often referred to as hard-wiring neuron paths. A Dropout layer, placed within a neural network, is implemented as a function operating during the training process and randomly discarding a specified fraction p of previous layer outputs and multiplying the remaining values by $1/(1-p)$. This forces the pattern recognition ability to be dispersed across the network, as during evaluation of every training step, a random part of the network is acting as if it was not gone. Once the training is completed, all Dropout layers are deactivated and simply pass all information forward, so that the model returns to its deterministic character.

In the experiments performed in the present work, as later discussed in 2.5, both Dropout and weight decay were used to mitigate overfitting, with good effects shown in particular in Fig. 16.

Methods for avoiding overfitting typically come with one or more “hyperparameters” (i.e. parameters which control the training process) that can represent how much confidence is given to the training data versus prior knowledge. For instance, if a regularizer is used, the strength of the regularizer, λ , would be a hyperparameter. In the terms of the present work, it generally corresponds to how many features in the material descriptor can be considered non-essential to making predictions and therefore discarded systematically throughout the training. Furthermore, when using Dropout, the probability p is also a hyperparameter.

One typically trains the model on the training dataset using a number of different hyperparameters and then subsequently chooses the best set of them using the validation error. This allows the determination of hyperparameter values that are appropriate to the problem at hand. However, in order to ensure that the determined hyperparameter values are not overly specific to the validation set, the final accuracy of the model is evaluated on a test set that was not used at all during training [46].




Fig. 17. A conceptual drawing depicting how overfitting mitigation effort can improve performance beyond regions with high known data density.

An additional advantage of mitigating overfitting to known data can be increased performance during extrapolation, as depicted conceptually in Fig. 17. This is thanks to reduced model complexity, which forces recognition of stronger and more broadly exhibited patterns rather than small deviations present in the training data, whether real or due to noise, that can significantly degrade the extrapolation capability of the ML model. It is important to recognize that cost of such model simplification is often reduced performance on previously unseen data that lays within the known region.

A.3. Transfer learning

Finally, one should consider the technique of transfer learning, which has been observed among deep learning models across a variety of domains [56–59]. Transfer learning refers to the ability of properly trained deep learning models to ‘transfer’ their knowledge to related tasks. In the least complex approach, one does this by simply ‘fine-tuning’ the parameters of the model using new training data (from the new task). This has to be done using a small learning rate and a small number of iterations on a loss function defined by the new training data. It has been observed that this often produces accurate results on the new task for a relatively small amount of additional data.

As an illustrative example, in [57], a network is first trained to recognize lower case handwritten characters. It is then shown that with minimal ‘fine-tuning,’ such a network can be made to accurately recognize upper case characters. The same phenomenon was also observed with a network that was first trained to recognize Chinese characters. Considering that this behavior has been widely observed [56,58,59], this shows that deep neural networks are often able to transfer knowledge between different but related tasks.

the present work adds to this evidence by showing that a network trained on the knowledge from the OQMD database covering a broad yet limited spectrum of material, can be easily adjusted to materials outside this spectrum with very little cost relative to the initial training. Specifically, the set of all (243) Fe–Ni–Cr σ -phase endmembers, described in 2.4, is shown in 3.5 to require transfer of only a few examples from that set to dramatically improve model performance on the rest.

Appendix B. Intermediate neural network models

The neural network design process was conducted in incremental fashion, starting from a perceptron, which is the simplest type of neural network proposed by Frank Rosenblatt in 1957 [82]. It effectively

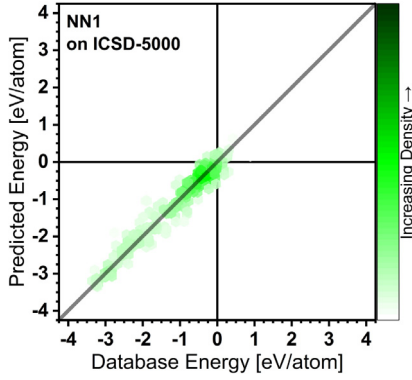


Fig. 18. Test of perceptron trained on the data from the first 5000 entries in the ICSD dataset and evaluated on the test set of 230 randomly selected entries ($\approx 5\%$).

operates as a linear function $f(\vec{d}) = A(w_1 d_1 + w_2 d_2 + \dots + w_n d_n)$ where d_i is i th element of the descriptor \vec{d} , w_i is the weight associated with it, and A is an activation function that can introduce non-linearity or turn it into a classifier. Here, the popular Sigmoid activation function was used.

The perceptron was first trained on the data from the first 5000 entries in the ICSD, to check whether the training was set up correctly. It achieved a MAE of 195 meV/atom on the test set of 230 randomly selected entries ($\approx 5\%$ from 5000). Results are shown in Fig. 18. When trained on the data from all entries in the ICSD, it achieved an MAE of 364 meV/atom on the test set ($\approx 5\%$ from 32116). This error is comparable to the performance of a random forest model based on PRDF (370 meV/atom), is slightly worse than a CM (250 meV/atom), and is significantly worse than a random-forest model trained on the same descriptor (90 meV/atom), as reported by Ward et al. [31]. Part of the significance of these results is the evident quality of the descriptor, as the model achieved performance that would be considered excellent just a few years prior to the present work while being much less complex and computationally costly. Furthermore, it is important to note the time- and space-complexity of the perceptron model. Training the final network took less than 8 s compared to around 10,000 s reported for the aforementioned random-forest methods, and the resulting model occupied less than 1 kb of memory. Following the testing of a perceptron, which allowed rough estimation of the good size of the network (i.e. number of weights), the design of the actual architecture began. All of these steps are schematically depicted in 19.

Next, in a few steps, the size of the network was incrementally increased. First, a layer of 1000 neurons was introduced. This reduced the performance on the first 5000 entries in the ICSD, likely due to overfitting issues, as the data was very limited. Performance on the ICSD was improved, reducing the test MAE to 305 meV/atom on the test set, however. The introduction of the next two 1000-width layers further reduced the MAE to 215 meV/atom. Based on these results, it was estimated that introducing 4 hidden layers with Sigmoid activation function and widths of 10000, 10000, 1000, and 100 would provide good results when trained on the much larger OQMD.

After switching to OQMD, the network exhibited issues with convergence, often predicting a single value for all of the entries. To mitigate this, the descriptor (i.e. network input) was normalized by dividing every element by its maximum value across the whole dataset. This solved the issue. Next, to improve the training behavior, the activation functions were changed from only the Sigmoid function to a mix of Soft Sign, Exponential Linear Unit, and Sigmoid, which was found to work well. These steps improved both the predictive performance and reduced the time required to converge. The network architecture resulting from these steps (internally designated NN8/Simple Good Network in 19) was the first to improve performance compared to the Ward et al. approach [31], achieving an MAE of 42 meV/atom on the

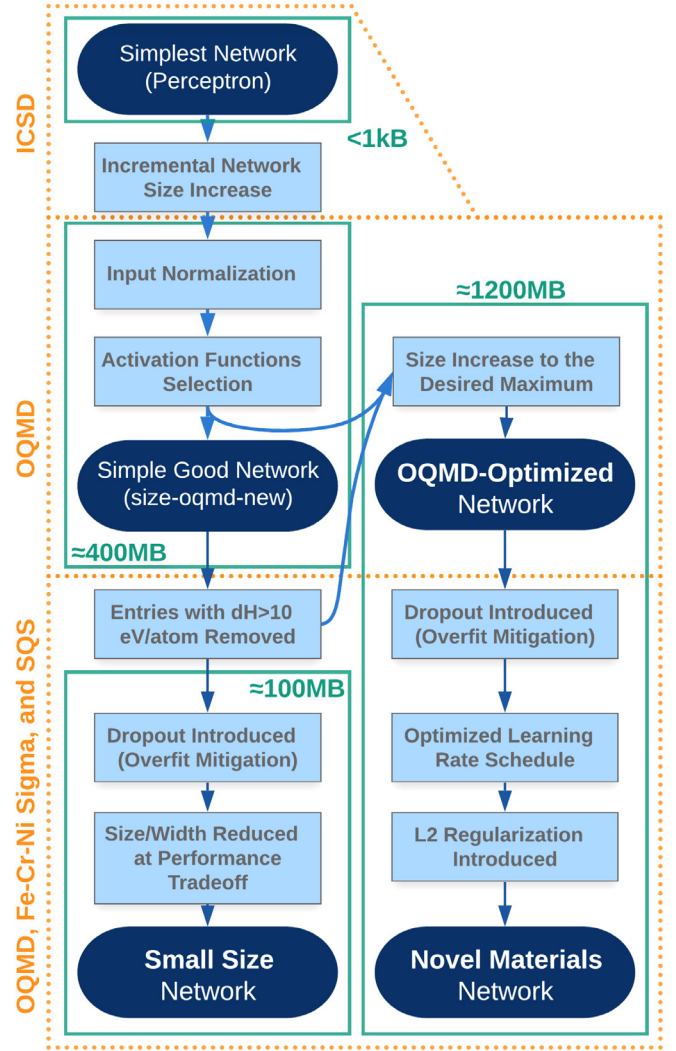


Fig. 19. The network design process schematic leading to the three final models.

test set of random subset 5% of the OQMD dataset. When testing this network, a small fraction of around 0.03% of likely incorrect entries in the OQMD was found, as described in 2.4, and was removed from the dataset used later in the design process.

Once a network with desired performance was obtained, the network size was increased until it either exceeded 1 GB or showed signs of extensive overfitting. At the first step of this process, two layers of width 10,000 were added, resulting in a network size of 1.2 GB and reduced overfitting, as indicated by the ratio of validation-to-training error lowered from 2.2 to 1.6, relative to NN8. The resulting network (internally designated NN9/OQMD-Optimized Network in 19), achieved an MAE of 28 meV/atom on the test set of random subset 5% of OQMD, which was the best performance on OQMD out of all the networks created in this project. If the 0.03% of abnormal data was not removed as described in 2.4, it would correspond to, on average, 6 data points which in one tested instance increased the MAE to 35 meV/atom. Important to point out, the training of this network was prone to staying in local minima at the beginning. The reported instance of the trained network exhibited no training progress between around rounds 5 and 25, after which its performance quickly increased. Detailed analysis of the performance is given in 3.2.

Once the main objective of the design process was obtained, i.e. the performance on the OQMD has improved appreciably beyond existing methods, the design process was focused on creating a tool for modeling




Fig. 20. The learning rate schedule used for training of more complex networks in the later stage of the design process (e.g., NN18).

materials that were not reported in the OQMD. Therefore, the objective changed from achieving the lowest MAE on a random subset 5% of OQMD to (1) reducing the mismatch between training and validation sets errors (i.e. difference between training accuracy and validation accuracy) during the training process, (2) keeping the test MAE on the OQMD below 50 meV/atom, and (3) improving performance on two material groups significantly different from the OQMD data, namely Special Quasirandom Structures (SQS) and Fe–Cr–Ni σ -phase (see 2.4).

With these new objectives, two Dropout layers in the middle part of the network were introduced to promote the distribution of pattern recognition abilities across the network. [83] This introduced a problem with convergence as the network became more likely to fall into local minima at the initial stages of the training, which was solved by introducing custom learning rate schedules. Specifically, the learning rate was initially set to a value orders of magnitude lower than during the default initial training and then ramped up to the previous (ADAM default setting in the majority of frameworks) learning rate of 0.001 (or above) after around 2 rounds of training. This type of learning rate schedule is known as a warm-up in the deep learning literature [84]. The schedule found to perform the best is presented in Fig. 20.

The next step was the introduction of ℓ^2 regularization, which is a technique that favors simplification of the descriptor and effectively rejects features of the descriptor that do not contribute to prediction performance [55]. An overview on it is given in Section 2.2. In the models reported in the present work an ℓ^2 value of 10^{-6} was used. Higher values were found to stop the training at early stages, impairing the pattern recognition, or in extreme cases (above 10^{-3}) force the network to discard the input completely, resulting in constant or near-constant output (i.e. mean value from the training dataset predicted for any structure).

The final step was small curation of the training data based on the OQMD-reported structure stability, i.e. the energy difference between the formation energy and the energy convex hull. The motivation for that was the notion that DFT results are inherently less accurate for unstable phases. In this step, all entries with energies of more than 2000 meV/atom above the convex hull were removed from the training set. Importantly, the validation and testing sets were not modified for consistent performance reporting.

All of these changes resulted in a neural network that has been optimized for predicting new materials. In the code and Supplementary materials, it is designated as NN20 (Novel Materials Network in 19). Compared to the OQMD-optimized network it was derived from, the test MAE on the OQMD increased from 28 to 49 meV/atom. However, at the same time, the mismatch between the training and validation set was reduced from 1.57 to 1.38. Or, as presented earlier in Fig. 16, reduced to about 1.15 for the same training duration. Furthermore, a relatively large portion of this error can be attributed to some unstable structures that were removed from the training set, but not from the test set. Once entries with formation energies of more than 1000 meV/atom above the convex hull were removed, the test MAE decreased to only

38 meV/atom. Restricting the test set further to only somewhat stable structures (stability below 250 meV/atom) resulted in an MAE of 30 meV/atom.

While the new-material-optimized network presented an increased MAE across a random subset of the OQMD, performance has significantly improved on the Fe–Cr–Ni σ -phase described in 2.4. The MAE has decreased from 55 to 41 meV/atom, indicating that the model based on this neural network is more capable of making predictions for new materials.

Once two performance-oriented models were developed, increasing the performance-to-cost ratio has been explored with the motivation that some studies would benefit from many times higher throughput at minor accuracy decrease. Architecture design started from the selection of a network with a balanced size-to-performance ratio (NN8) and the introduction of an overfitting mitigation technique (Dropout [47]) used for the network optimized for new materials, as depicted in Section 2.5. Next, the network was gradually narrowed (fewer neurons in layers) until the performance started to noticeably deteriorate (41.9 meV/atom for 5000- and 4000-width vs 42.1 for 3000-width). This approach allowed a significant reduction of the network size (and the computational intensity to run it) from around 1200 MB of the two other models to around 145 MB. If an application demands even more of a reduction in model size and computational cost, the same procedure could be continued until some minimum required performance is retained.

Appendix C. Feature ranking

Descriptor feature	Normalized squared weights sum
mean_NeighDiff_shell1_MeltingT	1
mean_MeltingT	0.97502
max_MeltingT	0.73512
mean_NeighDiff_shell1_NdUnfilled	0.69157
MaxPackingEfficiency	0.68889
most_MeltingT	0.67373
dev_GSvolume_pa	0.61042
var_NeighDiff_shell1_Columnn	0.58782
var_NeighDiff_shell1_CovalentRadius	0.57826
var_NeighDiff_shell1_MeltingT	0.57259
maxdiff_GSvolume_pa	0.55156
dev_MeltingT	0.5286
mean_SpaceGroupNumber	0.51761
min_MeltingT	0.50437
var_CellVolume	0.49467
var_NeighDiff_shell1_MendeleeevNumber	0.492
min_NeighDiff_shell1_MeltingT	0.47853
mean_NeighDiff_shell1_Columnn	0.45566
maxdiff_CovalentRadius	0.42998
var_NeighDiff_shell1_Electronegativity	0.42642
var_EffectiveCoordination	0.40506
min_NeighDiff_shell1_Columnn	0.39822
dev_NdUnfilled	0.39739
dev_CovalentRadius	0.36935
range_NeighDiff_shell1_Columnn	0.35956
range_NeighDiff_shell1_CovalentRadius	0.34585
mean_WCMagnitude_Shell1	0.34275

Descriptor feature	Normalized squared weights sum
mean_NeighDiff_shell1_MendelevNumber	0.33911
mean_EffectiveCoordination	0.33899
mean_Number	0.33769
mean_NdUnfilled	0.33408
maxdiff_MeltingT	0.33348
mean_AtomicWeight	0.33149
mean_NeighDiff_shell1_NdValence	0.33142
range_NeighDiff_shell1_MeltingT	0.33107
max_NfUnfilled	0.33041
dev_Electronegativity	0.33001
mean_NeighDiff_shell1_CovalentRadius	0.32999
var_NeighDiff_shell1_NdUnfilled	0.31973
dev_Column	0.31662
var_NeighDiff_shell1_NdValence	0.31481
mean_WCMagnitude_Shell2	0.31359
most_NfUnfilled	0.30916
MeanIonicChar	0.30732
mean_NeighDiff_shell1_Electronegativity	0.30277
min_EffectiveCoordination	0.29705
min_NeighDiff_shell1_CovalentRadius	0.29392
max_NeighDiff_shell1_GSvolume_pa	0.2875
most_SpaceGroupNumber	0.28472
max_NdUnfilled	0.28424
maxdiff_NdUnfilled	0.28405
var_NeighDiff_shell1_GSvolume_pa	0.28008
min_BondLengthVariation	0.27922
var_MeanBondLength	0.2768
dev_NdValence	0.27566
max_NeighDiff_shell1_MeltingT	0.27097
max_BondLengthVariation	0.26565
mean_NfValence	0.26558
mean_NsUnfilled	0.2612
max_NeighDiff_shell1_CovalentRadius	0.26026
max_GSvolume_pa	0.25985
min_GSvolume_pa	0.25895
mean_NdValence	0.25573
mean_NeighDiff_shell1_GSvolume_pa	0.25299
max_NValence	0.24749
range_NeighDiff_shell1_NdUnfilled	0.24643
max_CovalentRadius	0.23136
CanFormIonic	0.23135
min_NeighDiff_shell1_Electronegativity	0.22873
min_SpaceGroupNumber	0.22766
max_Electronegativity	0.22609
max_NdValence	0.22576
most_NdUnfilled	0.22198
min_NeighDiff_shell1_MendelevNumber	0.21991
var_NeighDiff_shell1_NpValence	0.21609
min_NeighDiff_shell1_NdUnfilled	0.2114
dev_SpaceGroupNumber	0.2099
most_NfValence	0.20888
min_MeanBondLength	0.2086
mean_BondLengthVariation	0.20507
var_NeighDiff_shell1_Row	0.20454
max_NeighDiff_shell1_NdUnfilled	0.20318
min_NeighDiff_shell1_NdValence	0.20123
min_CovalentRadius	0.19974
range_NeighDiff_shell1_MendelevNumber	0.19591
min_NeighDiff_shell1_GSvolume_pa	0.19565
most_NpUnfilled	0.19457
maxdiff_NUnfilled	0.19316
max_NeighDiff_shell1_NdValence	0.19307

Descriptor feature	Normalized squared weights sum
max_NpValence	0.1929
range_NeighDiff_shell1_GSvolume_pa	0.19166
most_NdValence	0.1904
max_MeanBondLength	0.19021
maxdiff_NfUnfilled	0.18897
max_NeighDiff_shell1_Column	0.18518
range_NeighDiff_shell1_Electronegativity	0.18322
var_NeighDiff_shell1_SpaceGroupNumber	0.18313
dev_NpValence	0.18099
mean_NpUnfilled	0.18091
range_NeighDiff_shell1_SpaceGroupNumber	0.17858
dev_MendelevNumber	0.17753
MaxIonicChar	0.176
mean_Column	0.17206
min_Electronegativity	0.17164
mean_WCMagnitude_Shell3	0.17077
mean_Row	0.17035
min_NeighDiff_shell1_SpaceGroupNumber	0.17031
most_NsUnfilled	0.16714
var_BondLengthVariation	0.16653
var_NeighDiff_shell1_NfUnfilled	0.16223
range_NeighDiff_shell1_NdValence	0.16094
frac_fValence	0.1609
maxdiff_Column	0.16083
max_NUnfilled	0.15916
mean_NpValence	0.15639
maxdiff_NpValence	0.15637
mean_MendelevNumber	0.15491
most_Electronegativity	0.15469
mean_Electronegativity	0.15458
max_SpaceGroupNumber	0.15429
dev_Row	0.15382
maxdiff_MendelevNumber	0.15373
var_NeighDiff_shell1_NpUnfilled	0.15135
max_NeighDiff_shell1_Electronegativity	0.15115
most_NUnfilled	0.14955
max_GSbandgap	0.14945
mean_NeighDiff_shell1_NUnfilled	0.14891
maxdiff_NValence	0.14819
mean_NeighDiff_shell1_NpValence	0.14768
maxdiff_NdValence	0.14735
max_NpUnfilled	0.14647
maxdiff_Electronegativity	0.14523
min_MendelevNumber	0.14119
mean_CovalentRadius	0.14049
mean_NeighDiff_shell1_Row	0.13945
maxdiff_GSbandgap	0.13891
max_NeighDiff_shell1_MendelevNumber	0.13858
most_Number	0.13823
most_AtomicWeight	0.13798
max_NeighDiff_shell1_NpValence	0.13757
Comp_L10Norm	0.13598
min_Row	0.13596
range_NeighDiff_shell1_NpValence	0.13524
mean_GSvolume_pa	0.1331
max_NeighDiff_shell1_NUnfilled	0.13205
mean_NeighDiff_shell1_NfValence	0.12888
min_NeighDiff_shell1_NpUnfilled	0.12778
mean_NeighDiff_shell1_SpaceGroupNumber	0.12722
mean_NsValence	0.12642
most_CovalentRadius	0.12616
var_NeighDiff_shell1_NUnfilled	0.12525

Descriptor feature	Normalized squared weights sum
mean_NeighDiff_shell1_Number	0.12466
Comp_L7Norm	0.12293
mean_NeighDiff_shell1_AtomicWeight	0.12229
min_NeighDiff_shell1_NpValence	0.12026
max_EffectiveCoordination	0.11995
min_NdValence	0.11984
maxdiff_NpUnfilled	0.11976
mean_NeighDiff_shell1_NsUnfilled	0.11836
max_NeighDiff_shell1_GSbandgap	0.11657
min_NUnfilled	0.11648
most_Column	0.1164
var_NeighDiff_shell1_Number	0.11483
most_MendeleevNumber	0.11312
max_NeighDiff_shell1_SpaceGroupNumber	0.11292
var_NeighDiff_shell1_AtomicWeight	0.11234
most_NpValence	0.11231
frac_dValence	0.11126
NComp	0.11097
min_Number	0.11062
range_NeighDiff_shell1_NpUnfilled	0.11002
dev_NValance	0.10868
min_Column	0.10846
max_NeighDiff_shell1_NpUnfilled	0.10837
maxdiff_Row	0.10735
Comp_L5Norm	0.10726
mean_NeighDiff_shell1_NpUnfilled	0.10682
maxdiff_SpaceGroupNumber	0.10604
dev_GSbandgap	0.10604
max_AtomicWeight	0.10495
max_GSmagmom	0.10416
maxdiff_GSmagmom	0.1039
dev_NUnfilled	0.10336
var_NeighDiff_shell1_NfValence	0.10059
dev_GSmagmom	0.10046
most_GSbandgap	0.09997
var_NeighDiff_shell1_NValance	0.09842
min_NeighDiff_shell1_Row	0.09798
min_NeighDiff_shell1_NUnfilled	0.09563
most_Row	0.09538
max_Number	0.0925
most_GSvolume_pa	0.09166
mean_GSbandgap	0.09097
range_NeighDiff_shell1_Row	0.09081
mean_NValance	0.0889
mean_NeighDiff_shell1_NsValence	0.08449
min_NsValence	0.08408
frac_pValence	0.08403
mean_NUnfilled	0.08244
mean_NfUnfilled	0.08194
dev_NpUnfilled	0.0818
dev_Number	0.08065
max_NeighDiff_shell1_GSmagmom	0.08049
max_Column	0.07989
min_AtomicWeight	0.07959
Comp_L3Norm	0.07913
max_NeighDiff_shell1_Row	0.0776
mean_NeighDiff_shell1_NValance	0.07619
mean_NeighDiff_shell1_NfUnfilled	0.07413
range_NeighDiff_shell1_NfUnfilled	0.07381
min_NValance	0.07297
max_NeighDiff_shell1_NValance	0.0726
range_NeighDiff_shell1_NfValence	0.07163

Descriptor feature	Normalized squared weights sum
min_NdUnfilled	0.07145
most_NsValence	0.07114
mean_NeighDiff_shell1_GSbandgap	0.06709
max_NfValence	0.06661
dev_AtomicWeight	0.06581
maxdiff_Number	0.06576
max_NeighDiff_shell1_NfUnfilled	0.06523
dev_NfUnfilled	0.06477
dev_NfValence	0.06373
range_NeighDiff_shell1_GSmagmom	0.06305
var_NeighDiff_shell1_NsUnfilled	0.06288
min_NeighDiff_shell1_Number	0.0623
frac_sValence	0.06099
min_NeighDiff_shell1_NfValence	0.06033
max_Row	0.05998
min_NeighDiff_shell1_NValance	0.05844
range_NeighDiff_shell1_NUnfilled	0.05819
var_NeighDiff_shell1_GSbandgap	0.05683
range_NeighDiff_shell1_AtomicWeight	0.0568
Comp_L2Norm	0.05638
min_NeighDiff_shell1_NsUnfilled	0.05541
most_NValance	0.0553
maxdiff_NsValence	0.05459
range_NeighDiff_shell1_NValance	0.0537
min_NeighDiff_shell1_AtomicWeight	0.05369
max_NsValence	0.05329
range_NeighDiff_shell1_GSbandgap	0.05299
min_NeighDiff_shell1_NfUnfilled	0.05266
maxdiff_NfValence	0.05147
dev_NsUnfilled	0.04884
max_MendeleevNumber	0.04844
maxdiff_AtomicWeight	0.04814
max_NeighDiff_shell1_NsUnfilled	0.04675
max_NeighDiff_shell1_NsValence	0.04663
var_NeighDiff_shell1_GSmagmom	0.04635
range_NeighDiff_shell1_Number	0.04416
max_NeighDiff_shell1_NfValence	0.04376
mean_NeighDiff_shell1_GSmagmom	0.0433
most_GSmagmom	0.04239
range_NeighDiff_shell1_NsUnfilled	0.03954
min_NeighDiff_shell1_NsValence	0.03932
max_NeighDiff_shell1_AtomicWeight	0.03905
max_NeighDiff_shell1_Number	0.03815
min_NfValence	0.03794
dev_NsValence	0.0373
maxdiff_NsUnfilled	0.03558
min_NfUnfilled	0.03537
min_NeighDiff_shell1_GSmagmom	0.03353
var_NeighDiff_shell1_NsValence	0.02948
min_NpValence	0.02946
max_NsUnfilled	0.02933
min_NeighDiff_shell1_GSbandgap	0.02735
mean_GSmagmom	0.02402
min_NpUnfilled	0.02233
range_NeighDiff_shell1_NsValence	0.02171
min_NsUnfilled	0.02051
min_GSbandgap	0.01299
min_GSmagmom	0.00132

Appendix D. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.commat.2022.111254>.

References

- [1] J.E. Saal, S. Kirklin, M. Aykol, B. Meredig, C. Wolverton, Materials design and discovery with high-throughput density functional theory: The open quantum materials database (OQMD), *JOM* 65 (2013) 1501–1509, <http://dx.doi.org/10.1007/s11837-013-0755-4>, <http://link.springer.com/10.1007/s11837-013-0755-4>.
- [2] S. Kirklin, J.E. Saal, B. Meredig, A. Thompson, J.W. Doak, M. Aykol, S. Rühl, C. Wolverton, The open quantum materials database (OQMD): Assessing the accuracy of DFT formation energies, *NPJ Comput. Mater.* 1 (2015) 15010, <http://dx.doi.org/10.1038/npjcompumats.2015.10>, <http://www.oqmd.org/download>, <http://www.nature.com/articles/npjcompumats201510>.
- [3] A. van de Walle, C. Nataraj, Z.K. Liu, The thermodynamic database database, CALPHAD, *Comput. Coupling Phase Diagr. Thermochem.* 61 (2018) 173–178, <http://dx.doi.org/10.1016/j.calphad.2018.04.003>.
- [4] A. Jain, S.P. Ong, G. Hautier, W. Chen, W.D. Richards, S. Dacek, S. Cholia, D. Gunter, D. Skinner, G. Ceder, K.A. Persson, Commentary: The materials project: A materials genome approach to accelerating materials innovation, *APL Mater.* 1 (2013) 011002, <http://dx.doi.org/10.1063/1.4812323>, <http://aip.scitation.org/doi/10.1063/1.4812323>.
- [5] S. Curtarolo, W. Setyawan, G.L.W. Hart, M. Jahnatek, R.V. Chepulskii, R.H. Taylor, S. Wang, J. Xue, K. Yang, O. Levy, M.J. Mehl, H.T. Stokes, D.O. Demchenko, D. Morgan, AFLOW: An automatic framework for high-throughput materials discovery, 2013, <http://dx.doi.org/10.1016/j.commat.2012.02.005>, www.aflowlib.org, [arXiv:abs/1308.5715](http://arxiv.org/abs/1308.5715).
- [6] C. Toher, C. Oses, D. Hicks, E. Gossett, F. Rose, P. Nath, D. Usanmaz, D.C. Ford, E. Perim, C.E. Calderon, J.J. Plata, Y. Lederer, M. Jahnatek, W. Setyawan, S. Wang, J. Xue, K. Rasch, R.V. Chepulskii, R.H. Taylor, G. Gomez, H. Shi, A.R. Supka, R.A.R. Al Arabi, P. Gopal, F.T. Cerasoli, L. Liyanage, H. Wang, I. Siloi, L.A. Agapito, C. Nyshadham, G.L.W. Hart, J. Carrete, F. Legrain, N. Mingo, E. Zurek, O. Isayev, A. Tropsha, S. Sanvito, R.M. Hanson, I. Takeuchi, M.J. Mehl, A.N. Kolmogorov, K. Yang, P. D'Amico, A. Calzolari, M. Costa, R.D. Gennaro, M.B. Nardelli, M. Fornari, O. Levy, S. Curtarolo, The AFLOW fleet for materials discovery, in: *Handbook of Materials Modeling*, Springer International Publishing, Cham, 2018, pp. 1–28, http://dx.doi.org/10.1007/978-3-319-42913-7_63-1, http://link.springer.com/10.1007/978-3-319-42913-7_63-1.
- [7] G. Pizzi, A. Cepellotti, R. Sabatini, N. Marzari, B. Kozinsky, AiIDA: Automated interactive infrastructure and database for computational science, *Comput. Mater. Sci.* 111 (2016) 218–230, <http://dx.doi.org/10.1016/j.commat.2015.09.013>, <https://linkinghub.elsevier.com/retrieve/pii/S0927025615005820>.
- [8] O. Isayev, C. Oses, C. Toher, E. Gossett, S. Curtarolo, A. Tropsha, Universal fragment descriptors for predicting properties of inorganic crystals, *Nature Commun.* 8 (2017) 15679, <http://dx.doi.org/10.1038/ncomms15679>, <http://www.nature.com/articles/ncomms15679>.
- [9] F. Legrain, J. Carrete, A. Van Rooyeghem, S. Curtarolo, N. Mingo, How chemical composition alone can predict vibrational free energies and entropies of solids, *Chem. Mater.* 29 (2017) 6220–6227, <http://dx.doi.org/10.1021/acs.chemmater.7b00789>.
- [10] G. Pilania, X.-Y. Liu, Machine learning properties of binary wurtzite superlattices, *J. Mater. Sci.* 53 (1987) <http://dx.doi.org/10.1007/s10853-018-1987-z>.
- [11] J. Jung, J.I. Yoon, H.K. Park, J.Y. Kim, H.S. Kim, Bayesian approach in predicting mechanical properties of materials: Application to dual phase steels, *Mater. Sci. Eng. A* 743 (2019) 382–390, <http://dx.doi.org/10.1016/j.msea.2018.11.106>.
- [12] B. Ouyang, Y. Wang, Y. Sun, G. Ceder, Computational investigation of halogen-substituted Na argyrodites as solid-state superionic conductors, *Chem. Mater.* 32 (2020) 1896–1903, <http://dx.doi.org/10.1021/acs.chemmater.9b04541>.
- [13] B.J. Bucior, N.S. Bobbitt, T. Islamoglu, S. Goswami, A. Gopalan, T. Yildirim, O.K. Farha, N. Bagheri, R.Q. Snurr, Energy-based descriptors to rapidly predict hydrogen storage in metal-organic frameworks †, *Mol. Syst. Des. Eng.* 4 (2019) 162, <http://dx.doi.org/10.1039/c8me00050f>.
- [14] A. Chandrasekaran, D. Kamal, R. Batra, C. Kim, L. Chen, R. Ramprasad, Solving the electronic structure problem with machine learning, *Npj Comput. Mater.* 5 (2019) <http://dx.doi.org/10.1038/s41524-019-0162-7>.
- [15] K. Kim, L. Ward, J. He, A. Krishna, A. Agrawal, C. Wolverton, Machine-learning-accelerated high-throughput materials screening: Discovery of novel quaternary Heusler compounds machine learning accelerated high-throughput materials screening: Discovery of novel quaternary Heusler compounds, *Phys. Rev. Mater.* 2 (2018) 123801, <http://dx.doi.org/10.1103/PhysRevMaterials.2.123801>.
- [16] C. Wen, Y. Zhang, C. Wang, D. Xue, Y. Bai, S. Antonov, L. Dai, T. Lookman, Y. Su, Machine learning assisted design of high entropy alloys with desired property, *Acta Mater.* 170 (2019) 109–117, <http://dx.doi.org/10.1016/j.actamat.2019.03.010>, <https://linkinghub.elsevier.com/retrieve/pii/S1359645419301430>.
- [17] L. Scime, J. Beuth, Using machine learning to identify in-situ melt pool signatures indicative of flaw formation in a laser powder bed fusion additive manufacturing process, *Addit. Manuf.* 25 (2019) 151–165, <http://dx.doi.org/10.1016/j.addma.2018.11.010>.
- [18] J. Schmidt, M.R.G. Marques, S. Botti, M.A.L. Marques, Recent advances and applications of machine learning in solid-state materials science, *NPJ Comput. Mater.* 5 (2019) 83, <http://dx.doi.org/10.1038/s41524-019-0221-0>, <http://www.nature.com/articles/s41524-019-0221-0>.
- [19] R.K. Vasudevan, K. Choudhary, A. Mehta, R. Smith, G. Kusne, F. Tavazza, L. Vlcek, M. Ziatdinov, S.V. Kalinin, J. Hattrick-Simpers, Materials science in the artificial intelligence age: High-throughput library generation, machine learning, and a pathway from correlations to the underpinning physics, *MRS Commun.* 9 (2019) 821–838, <http://dx.doi.org/10.1557/mrc.2019.95>.
- [20] M. Aykol, V.I. Hegde, L. Hung, S. Suram, P. Herring, C. Wolverton, J.S. Hummelshøj, Network analysis of synthesizable materials discovery, *Nature Commun.* 10 (2019) <http://dx.doi.org/10.1038/s41467-019-10030-5>.
- [21] V. I. Hegde, M. Aykol, S. Kirklin, C. Wolverton, The phase stability network of all inorganic materials, *Sci. Adv.* 6 (2020) <http://dx.doi.org/10.1126/sciadv.aay5606>.
- [22] O. Kononova, H. Huo, T. He, Z. Rong, T. Botari, W. Sun, V. Tshitoyan, G. Ceder, Text-mined dataset of inorganic materials synthesis recipes, *Sci. Data* 6 (2019) <http://dx.doi.org/10.1038/s41597-019-0224-1>.
- [23] K. Alberi, M.B. Nardelli, A. Zakutayev, L. Mitas, S. Curtarolo, A. Jain, M. Fornari, N. Marzari, I. Takeuchi, M.L. Green, M. Kanatzidis, M.F. Toney, S. Butenko, B. Meredig, S. Lany, U. Kattner, A. Davydov, E.S. Toberer, V. Stevanovic, A. Walsh, N.-G. Park, A. Aspuru-Guzik, D.P. Tabor, J. Nelson, J. Murphy, A. Setlur, J. Gregoire, H. Li, R. Xiao, A. Ludwig, L.W. Martin, A.M. Rappe, S.-H. Wei, J. Perkins, The 2019 materials by design roadmap, *J. Phys. D: Appl. Phys.* 52 (2019) 013001, <http://dx.doi.org/10.1088/1361-6463/AA926>.
- [24] P.V. Balachandran, A.A. Emery, J.E. Gubernatis, T. Lookman, C. Wolverton, A. Zunger, Predictions of new perovskite compounds by combining machine learning and density functional theory, *Phys. Rev. Mater.* 2 (2018) 43802, <http://dx.doi.org/10.1103/PhysRevMaterials.2.043802>.
- [25] Z. Li, Q. Xu, Q. Sun, Z. Hou, W.-J. Yin, Thermodynamic stability landscape of Halide double perovskites via high-throughput computing and machine learning, *Adv. Funct. Mater.* 29 (2019) 1807280, <http://dx.doi.org/10.1002/adfm.201807280>, <http://doi.wiley.com/10.1002/adfm.201807280>.
- [26] S. Im, S.L. Shang, N.D. Smith, A.M. Krajewski, T. Lichtenstein, H. Sun, B.J. Bocklund, Z.K. Liu, H. Kim, Thermodynamic properties of the Nd-Bi system via emf measurements, DFT calculations, machine learning, and CALPHAD modeling, *Acta Mater.* 223 (2022) 117448, <http://dx.doi.org/10.1016/j.actamat.2021.117448>.
- [27] S.-L. Shang, H. Sun, B. Pan, Y. Wang, A.M. Krajewski, M. Banu, J. Li, Z.-K. Liu, Forming Mechanism of Equilibrium and Non-equilibrium Metallurgical Phases in Dissimilar Materials: Illustrated With Aluminum/steel (Al-Fe) Joints, (2021). <http://dx.doi.org/10.21203/rs.3.rs-745143/v1>.
- [28] J.R. Hattrick-Simpers, K. Choudhary, C. Corgnale, A simple constrained machine learning model for predicting high-pressure-hydrogen-compressor materials, *Mol. Syst. Des. Eng.* 3 (2018) 509, <http://dx.doi.org/10.1039/c8me00005k>.
- [29] D. Jha, L. Ward, A. Paul, W.-k. Liao, A. Choudhary, C. Wolverton, A. Agrawal, ElemNet: Deep learning the chemistry of materials from only elemental composition, *Sci. Rep.* 8 (2018) 17593, <http://dx.doi.org/10.1038/s41598-018-35934-y>, <http://www.nature.com/articles/s41598-018-35934-y>.
- [30] L. Ward, A. Agrawal, A. Choudhary, C. Wolverton, A general-purpose machine learning framework for predicting properties of inorganic materials, *NPJ Comput. Mater.* 2 (2016) <http://dx.doi.org/10.1038/npjcompumats.2016.28>.
- [31] L. Ward, R. Liu, A. Krishna, V.I. Hegde, A. Agrawal, A. Choudhary, C. Wolverton, Including crystal structure attributes in machine learning models of formation energies via voronoi tessellations, *Phys. Rev. B* 96 (2017) 024104, <http://dx.doi.org/10.1103/PhysRevB.96.024104>, <http://link.aps.org/doi/10.1103/PhysRevB.96.024104>.
- [32] A. Seko, H. Hayashi, K. Nakayama, A. Takahashi, I. Tanaka, Representation of compounds for machine-learning prediction of physical properties, *Phys. Rev. B* 95 (2017) <http://dx.doi.org/10.1103/PhysRevB.95.144110>, <https://journals.aps.org/prb/abstract/10.1103/PhysRevB.95.144110>.
- [33] K.T. Schütt, H. Glawe, F. Brockherde, A. Sanna, K.R. Müller, E.K.U. Gross, How to represent crystal structures for machine learning: Towards fast prediction of electronic properties, *Phys. Rev. B* 89 (2014) 205118, <http://dx.doi.org/10.1103/PhysRevB.89.205118>, <https://link.aps.org/doi/10.1103/PhysRevB.89.205118>.
- [34] R.E.A. Goodall, A.A. Lee, Predicting materials properties without crystal structure: Deep representation learning from stoichiometry, *Nature Commun.* 11 (2020) 6280, <http://dx.doi.org/10.1038/s41467-020-19964-7>, <http://www.nature.com/articles/s41467-020-19964-7>.
- [35] L. Kaufman, H. Bernstein, Computer calculation of phase diagrams. With special reference to refractory metals, 1970, <http://inis.iaea.org/Search/search.aspx?orig.q=RN:2004171>.
- [36] Z.K. Liu, Ocean of data: Integrating first-principles calculations and CALPHAD modeling with machine learning, *J. Phase Equilibria Diffus.* 39 (2018) 635–649, <http://dx.doi.org/10.1007/S11669-018-0654-Z/FIGURES/8>, <https://link.springer.com/article/10.1007/s11669-018-0654-z>.

- [37] Z.-K. Liu, Y. Wang, S.-L. Shang, Zentropy theory for positive and negative thermal expansions, 2021, [arXiv:2107.06455v6](https://arxiv.org/abs/2107.06455v6).
- [38] A. Debnath, A.M. Krajewski, H. Sun, S. Lin, M. Ahn, W. Li, S. Priya, J. Singh, S. Shang, A.M. Beese, Z.-K. Liu, W.F. Reinhart, Generative deep learning as a tool for inverse design of high entropy refractory alloys, *J. Mater. Inform.* 1 (2021) 3, <http://dx.doi.org/10.20517/JMI.2021.05>, <https://jmijournal.com/article/view/4294>.
- [39] D. Jha, L. Ward, Z. Yang, C. Wolverton, I. Foster, W.-K. Liao, A. Choudhary, A. Agrawal, IRNet, in: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, ACM, New York, NY, USA, 2019, pp. 2385–2393, <http://dx.doi.org/10.1145/3292500.3330703>, <https://dl.acm.org/doi/10.1145/3292500.3330703>.
- [40] K.T. Schütt, H.E. Sauceda, P.-J. Kindermans, A. Tkatchenko, K.-R. Müller, SchNet – A Deep learning architecture for molecules and materials, *J. Chem. Phys.* 148 (2018) 241722, <http://dx.doi.org/10.1063/1.5019779>, <http://aip.scitation.org/doi/10.1063/1.5019779>.
- [41] FAIR Principles - GO FAIR, <https://www.go-fair.org/fair-principles/>.
- [42] C. Rycroft, Voro++: A Three-Dimensional Voronoi Cell Library In C++, Tech. Rep., Lawrence Berkeley National Lab.(LBNL), Berkeley, CA, United States, 2009.
- [43] V.N. Vapnik, An overview of statistical learning theory, *IEEE Trans. Neural Netw.* 10 (1999) 988–999.
- [44] I. Goodfellow, Y. Bengio, A. Courville, Deep Learning, MIT Press, 2016.
- [45] D.P. Kingma, J. Ba, Adam: A method for stochastic optimization, 2014, arXiv preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [46] T. Hastie, R. Tibshirani, J. Friedman, The Elements of Statistical Learning: Data Mining, Inference, and Prediction, Springer Science & Business Media, 2009.
- [47] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, R. Salakhutdinov, Dropout: A simple way to prevent neural networks from overfitting, *J. Mach. Learn. Res.* 15 (2014) 1929–1958.
- [48] T. Chen, M. Li, Y. Li, M. Lin, N. Wang, M. Wang, U.W. Cmu, S. Nus, T. Nyu, T. Xiao, B. Xu, C. Zhang, Z. Zhang, M.U. Alberta, MXNet: A Flexible and Efficient Machine Learning Library for Heterogeneous Distributed Systems, Tech. Rep..
- [49] G. Kresse, J. Hafner, Ab initio molecular dynamics for liquid metals, *Phys. Rev. B* 47 (1993) 558–561, <http://dx.doi.org/10.1103/PhysRevB.47.558>, <https://link.aps.org/doi/10.1103/PhysRevB.47.558>.
- [50] M. Feurer, B. Bocklund, S. Shang, A. Beese, Z.-K. Liu, Cr-Fe-Ni Sigma Phase Finite Temperature Calculations, *Citrine Inform.*, 2019, <http://dx.doi.org/10.25920/YJRC-ZJ59>.
- [51] C.-C. Hsieh, W. Wu, Overview of intermetallic sigma (σ) phase precipitation in stainless steels, *ISRN Metall.* 2012 (2012) 1–16, <http://dx.doi.org/10.5402/2012/732471>, <https://www.hindawi.com/journals/isrn/2012/732471/>.
- [52] A. Zunger, S.-H. Wei, L.G. Ferreira, J.E. Bernard, Special Quasirandom Structures, *Tech. Rep.*, 65, 1990.
- [53] C. Jiang, C. Wolverton, J. Sofo, L.-Q. Chen, Z.-K. Liu, First-principles study of binary BCC alloys using special quasirandom structures, *Phys. Rev. B* 69 (2004) 214202, <http://dx.doi.org/10.1103/PhysRevB.69.214202>, <https://link.aps.org/doi/10.1103/PhysRevB.69.214202>.
- [54] D. Shin, R. Arróyave, Z.-K. Liu, A. Van de Walle, Thermodynamic properties of binary HCP solution phases from special quasirandom structures, *Phys. Rev. B* 74 (2006) 024204, <http://dx.doi.org/10.1103/PhysRevB.74.024204>, <https://link.aps.org/doi/10.1103/PhysRevB.74.024204>.
- [55] L2 Regularization, https://www.textbook.ds100.org/ch/16/reg_ridge.html.
- [56] C. Tan, F. Sun, T. Kong, W. Zhang, C. Yang, C. Liu, A survey on deep transfer learning, in: International Conference on Artificial Neural Networks, Springer, 2018, pp. 270–279.
- [57] D.C. Cireşan, U. Meier, J. Schmidhuber, Transfer learning for latin and Chinese characters with deep neural networks, in: The 2012 International Joint Conference on Neural Networks, IJCNN, IEEE, 2012, pp. 1–6.
- [58] H. Chang, J. Han, C. Zhong, A.M. Snijders, J.-H. Mao, Unsupervised transfer learning via multi-scale convolutional sparse coding for biomedical applications, *IEEE Trans. Pattern Anal. Mach. Intell.* 40 (2017) 1182–1194.
- [59] D. George, E. Huerta, Deep learning for real-time gravitational wave detection and parameter estimation: Results with advanced LIGO data, *Phys. Lett. B* 778 (2018) 64–70.
- [60] POSCAR file, https://cms.mpi.univie.ac.at/vasp/vasp/POSCAR_file.html.
- [61] S.R. Hall, F.H. Allen, I.D. Brown, The crystallographic information file (CIF): A new standard archive file for crystallography, *Acta Crystallogr. Sect. A* 47 (1991) 655–685, <http://dx.doi.org/10.1107/S010876739101067X>, <https://onlinelibrary.wiley.com/doi/10.1107/S010876739101067X>.
- [62] S.P. Ong, W.D. Richards, A. Jain, G. Hautier, M. Kocher, S. Cholia, D. Gunter, V.L. Chevrier, K.A. Persson, G. Ceder, Python materials genomics (pymatgen): A robust, open-source python library for materials analysis, *Comput. Mater. Sci.* 68 (2013) 314–319, <http://dx.doi.org/10.1016/j.commatsci.2012.10.028>, <https://linkinghub.elsevier.com/retrieve/pii/S0927025612006295>.
- [63] Y. LeCun, Y. Bengio, G. Hinton, Deep learning, *Nature* 521 (2015) 436–444.
- [64] D. Silver, A. Huang, C.J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, et al., Mastering the game of go with deep neural networks and tree search, *Nature* 529 (2016) 484.
- [65] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, Bert: Pre-training of deep bidirectional transformers for language understanding, 2018, arXiv preprint [arXiv:1810.04805](https://arxiv.org/abs/1810.04805).
- [66] G. Carleo, M. Troyer, Solving the quantum many-body problem with artificial neural networks, *Science* 355 (2017) 602–606.
- [67] F. Rosenblatt, The perceptron: A probabilistic model for information storage and organization in the brain, *Psychol. Rev.* 65 (1958) 386.
- [68] Y. LeCun, L. Jackel, L. Bottou, A. Brunot, C. Cortes, J. Denker, H. Drucker, I. Guyon, U. Muller, E. Sackinger, et al., Comparison of learning algorithms for handwritten digit recognition, in: International Conference on Artificial Neural Networks, vol. 60, Perth, Australia, 1995, pp. 53–60.
- [69] Y. LeCun, B.E. Boser, J.S. Denker, D. Henderson, R.E. Howard, W.E. Hubbard, L.D. Jackel, Handwritten digit recognition with a back-propagation network, in: Advances in Neural Information Processing Systems, 1990, pp. 396–404.
- [70] Y. LeCun, L. Bottou, Y. Bengio, P. Haffner, Gradient-based learning applied to document recognition, *Proc. IEEE* 86 (1998) 2278–2324.
- [71] A. Krizhevsky, I. Sutskever, G.E. Hinton, Imagenet classification with deep convolutional neural networks, in: Advances in Neural Information Processing Systems, 2012, pp. 1097–1105.
- [72] W.B. Charles Jr., A comparison of the accuracy of different functionals, *Chem. Phys. Lett.* 246 (1995) 40–44.
- [73] S. Alturk, D. Avci, O. Tamer, Y. Atalay, Comparison of different hybrid DFT methods on structural, spectroscopic, electronic and NLO parameters for a potential NLO material, *Comput. Theor. Chem.* 1100 (2017) 34–45, <http://dx.doi.org/10.1016/j.comptc.2016.12.007>, <https://linkinghub.elsevier.com/retrieve/pii/S2210271X16304935>.

- [74] V. Vapnik, *The Nature of Statistical Learning Theory*, Springer science & business media, 2013.
- [75] I.J. Goodfellow, Y. Bulatov, J. Ibarz, S. Arnoud, V. Shet, Multi-digit number recognition from street view imagery using deep convolutional neural networks, 2013, arXiv preprint [arXiv:1312.6082](https://arxiv.org/abs/1312.6082).
- [76] G.E. Dahl, D. Yu, L. Deng, A. Acero, Context-dependent pre-trained deep neural networks for large-vocabulary speech recognition, *IEEE/ACM Trans. Audio Speech Lang. Process.* 20 (2011) 30–42.
- [77] W. Huang, P. Martin, H.L. Zhuang, Machine-learning phase prediction of high-entropy alloys, *Acta Mater.* 169 (2019) 225–236, <http://dx.doi.org/10.1016/j.actamat.2019.03.012>, <https://linkinghub.elsevier.com/retrieve/pii/S1359645419301454>.
- [78] S. Feng, H. Zhou, H. Dong, Using deep neural network with small dataset to predict material defects, *Mater. Des.* 162 (2019) 300–310, <http://dx.doi.org/10.1016/j.matdes.2018.11.060>.
- [79] B. Everitt, A. Skrondal, *The Cambridge Dictionary of Statistics*, vol. 106, Cambridge University Press Cambridge, 2002.
- [80] A.N. Tikhonov, On the solution of ill-posed problems and the method of regularization, in: *Doklady Akademii Nauk*, vol. 151, Russian Academy of Sciences, 1963, pp. 501–504.
- [81] A.E. Hoerl, R.W. Kennard, Ridge regression: Biased estimation for nonorthogonal problems, *Technometrics* 12 (1970) 55–67.
- [82] F. Rosenblatt, *The Perceptron—A perceiving and recognizing automaton*, Report 85-460-1, 1957.
- [83] N. Srivastava, G. Hinton, A. Krizhevsky, R. Salakhutdinov, Dropout: A simple way to prevent neural networks from overfitting, *J. Mach. Learn. Res.* 15 (2014) 1929–1958.
- [84] A. Gotmare, N.S. Keskar, C. Xiong, R. Socher, A closer look at deep learning heuristics: Learning rate restarts, warmup and distillation, 2018, arXiv preprint [arXiv:1810.13243](https://arxiv.org/abs/1810.13243).