

# CrabNet for explainable deep learning in materials science: bridging the gap between academia and industry

Anthony Yu-Tung Wang · Mahamad Salah  
Mahmoud · Mathias Czasny · Aleksander  
Gurlo

Received: date  
Accepted: date

**Abstract** Despite recent breakthroughs in deep learning for materials informatics, there exists a disparity between their popularity in academic research and their limited adoption in the industry. A significant contributor to this “interpretability-adoption gap” is the prevalence of black-box models and the lack of built-in methods for model interpretation. While established methods for evaluating model performance exist, an intuitive understanding of the modeling and decision-making processes in models is nonetheless desired in many cases.

In this work, we demonstrate several ways of incorporating model interpretability to the structure-agnostic Compositionally Restricted Attention-Based network, CrabNet. We show that CrabNet learns meaningful, material property-specific element representations based solely on the data with no additional supervision. These element representations can then be used to explore element identity, similarity, behavior, and interactions within different chemical environments. Chemical compounds can also be uniquely represented and examined to reveal clear structures and trends within the chemical space. Additionally, visualizations of the attention mechanism can be used in conjunction to further understand the modeling process, identify potential modeling or dataset errors, and hint at further chemical insights leading to a better understanding of the phenomena governing material properties. We feel confident that the interpretability methods introduced in this work for CrabNet will be of keen interest to materials informatics researchers as well as industrial practitioners alike.

**Keywords** Materials Informatics · Deep Learning · Self-Attention · Interpretability · Explainable AI · XAI

---

A. Y.-T. Wang, M. Czasny, A. Gurlo  
Fachgebiet Keramische Werkstoffe/Chair of Advanced Ceramic Materials, Technische Universität Berlin,  
10623 Berlin, Germany  
E-mail: anthony.wang@ceramics.tu-berlin.de

M. S. Mahmoud  
Department of Computer and Data Sciences, Case Western Reserve University, Cleveland, Ohio 44106,  
USA

## Introduction

Machine learning (ML) in materials informatics (MI) has received significant attention in the academic research world and is gaining widespread adoption [1–5]. More specifically, it has recently been extensively studied for its use in the research and design of novel inorganic materials [6–10]. This is enabled by three major developments: (1) the increasing number of material property datasets as well as the improvement in dataset quality and variety, (2) the rapid pace and development of new ML models tailored to addressing different challenges in materials science (*e.g.*, regression, classification), supplemented by (3) the increase in available computing power and accessibility to ML and deep learning tools. The combination of these developments led to improved capabilities in the exploration and modeling of material properties in the academic world.

Classical ML methods (*e.g.*, linear regression, random forest, support vector machines) have successfully been used for the regression and classification of many material properties [11–17]. These methods usually rely on the featurization of the input chemical formulae into numerical features that are usable by the models. Typically, this is achieved through the use of a composition-based feature vector (CBFV), which uses descriptive statistics of the properties of constituent atoms in each compound to uniquely represent it [18]. Some common CBFV feature sets are Oliylyk, Magpie, Jarvis and mat2vec [11,12,19,20]. Here a distinction is made between physically-derived CBFVs (with features based on measurable element properties) like Oliylyk and Magpie and computationally-derived CBFVs (with features obtained from computational or deep learning models) like Jarvis and mat2vec. For some properties, additional features such as structural information, processing or measurement conditions are included to further improve model performance [2,21,22,16].

In more recent years, deep learning (DL) models have gained widespread popularity in MI due to numerous advantages compared to classical ML methods. Some examples are ElemNet, CGCNN, MEGNet, DimeNet++, and ALIGNN [23–27]. More recently, graph neural network (GNN) models incorporating attention-based mechanisms such as CrabNet, Roost and H-CLMP have gained increasing popularity [28–30]. GNNs have shown improved performance compared to other DL models, particularly in the absence of structural information as model inputs. Another advantage of GNNs is that the inductive biases built into the model and the input data structure are more suited to the learning of material properties, since the interactions between the atoms in the compound can be modeled as weighted interactions between nodes in a graph. In CrabNet, for example, the atom representations are either based on a CBFV feature (mat2vec element vectors) or a non-CBFV feature (onehot element vectors) [28]. For the sake of clarity, the remaining text will use the acronym DL to refer to both deep learning (DL) and graph neural network (GNN) models and methods.

Unfortunately, while DL methods show superb performance in modeling material properties, the element features used by these models typically do not represent any measurable physical property of the elements themselves. Instead, the element representations are learned from the data during the model training process. Therefore, they do not directly provide useful information or insights that can be interpreted by humans. This is different from the CBFV representation typically used in classical ML, where the features represent properties of the elements which are known *a priori*, such as the atomic mass, first ionization energy, or number of valence electrons.

Despite the high performance of the DL models, there is a disparity between their extensive study in academic research and their limited adoption in the industry for the exploration of materials. We term this disparity the “interpretability-adoption gap”. One significant hurdle to the widespread adoption of the often “black-box” models is the lack of built-in methods for model interpretation. While there are established methods of

evaluating model performance in academia [14, 31–33], those who are less familiar with DL typically require more intuition into how the models function before they can fully trust the results. Particularly in industry, where there is usually a lower risk tolerance compared to academia, findings based on black-box models and vague model evaluation criteria are not enough to justify making high-stakes decisions such as investing in new research [34–38, 5]. Tangible methods of investigating and understanding model decision-making processes are therefore required to facilitate their adoption in an industrial setting [39].

This led to the development of explainable AI (XAI), which aims to introduce methods for deciphering the internal workings of black-box models and thus enabling users to understand the modeling processes and results [39, 40]. Examples of XAI in research fields outside of MI include: visualizing word embeddings in Natural Language Processing [41–43], inspecting decision-making processes in reinforcement learning [44–46], visualizing pixel importances [47, 48], or segmenting in computer vision [49, 50]. To date, however, XAI techniques have—with the exception of a few works employing classical ML—largely been underexplored for DL in the MI field [51, 52, 10].

Two common post-hoc model-agnostic methods for obtaining explainable models in classical ML are SHAP and LIME [53, 54, 39, 55]. Both of these methods are built on top of existing black-box models and use local feature perturbation to estimate the contributions from input features towards the predictions. Other models such as random forest, gradient boosting, and lasso regression inherently provide model interpretability via the use of internal feature importance metrics and (in some models) through bootstrap sampling and feature sampling [56, 51, 39]. Nonetheless, these techniques require that the individual features of the input data are meaningful and represent a measurable feature or physical property. This works in the domain of classical ML and when using a physically-derived CBFV to featurize compounds; however, this is not the case for DL methods where the features typically do not reflect a measurable value. Thus, these traditional ways of model interpretability fall short in use for the DL models.

Therefore, it is the goal of this work to explore how to increase model interpretability in DL models specifically for applications in MI. Here, we demonstrate how parts of the typically black-box modeling process can be communicated visually and in an interpretable way, using our attention-based model, CrabNet [28]. We have extended CrabNet’s architecture to enable intrinsic interpretability using several methods to be discussed below. In this regard, we lay the first bricks in the bridge spanning the interpretability-adoption gap between academia and industry. This will not only aid researchers in further developing complex models with interpretability in focus, but also promote the adoption of these modeling methods in the materials science industry.

## Results & Discussion

The results of this study are described in five subsections. We first compare the element embeddings learned by CrabNet against other CBFV feature sets from the literature, and show how chemical behavior and patterns in element properties can be learned entirely from the training data for each material property. We also show that the learned element representations are comparable to physically-derived CBFVs. Secondly, as part of this analysis, we characterize the element prevalence imbalance in the datasets using the Shannon equitability index and relate that to the quality of the learned element embeddings. Third, we further examine how the element representations are successively updated using information about their chemical environment in the compounds, and how they may be used to gain additional insights about element behaviors in different environments. Fourth, we inspect how entire chemical compounds can also be adequately

captured using the EDMs and subsequently visualized. We identify interesting trends in the compound representations relating the bond character and number of elements in the compounds to the material property and prediction error, and discuss how such visualizations can lead to additional understanding about the modeling process and the underlying materials chemistry. Lastly, we explore how the self-attention mechanism in CrabNet can be visualized in the form of videos and used to further examine the modeling process, leading to potential new insights about the chemical interactions within a compound. While we use the OQMD\_Bandgap dataset to demonstrate the analyses, we note that similar analyses can be also carried out with any of the 28 materials datasets presented in this work.

### Learning meaningful and per-property elements representations

Element representations were obtained as featurized CBFVs, which are fixed-length vectors where each element is uniquely described by the same set of features [18, 12]. For the Oliyntyk, Magpie and mat2vec element property feature sets, we use the published vectors to represent the elements [18, 20]. For the CrabNet element representations, we extract the element vectors from the element-derived matrices (EDMs) at the output of the embedding layer (please refer to the CrabNet publication for architecture details [28]). We can examine the similarity between two element vectors  $x$  and  $y$  by computing the Pearson correlation coefficient  $r$  using Equation 1:

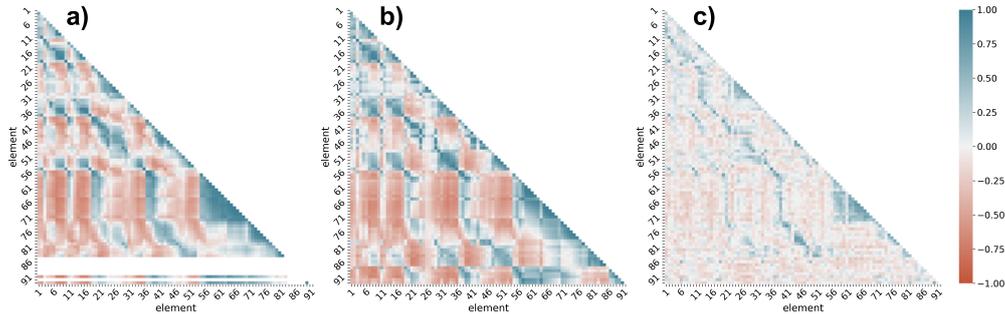
$$r = \frac{\sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^n (x_i - \bar{x})^2 (y_i - \bar{y})^2}} \quad (1)$$

where  $n$  is the number of features,  $x_i$  and  $y_i$  are the values of the  $i$ th feature, and  $\bar{x}$  and  $\bar{y}$  are the mean values of  $x$  and  $y$ , respectively.

The correlation  $r$  ranges from  $-1$  to  $1$ ; the higher or lower the value of  $r$  is, the more correlated or anticorrelated are the features that describe the elements, respectively. A value of zero means that there are no correlations between the features of the elements. We compute the pairwise correlation coefficients between the element vectors for all elements and for all element property representations, and show these as heatmaps in Figure 1. Note that the plots are cropped to the range of elements of the Oliyntyk heatmap to aid comparison; please refer to supplementary Figure S-1 in the supplementary information (SI) for the full heatmaps. In addition, interactive versions of the plots are provided in the SI.

Here we can observe that element vectors based on the Oliyntyk and Magpie CBFVs contain large regions of similar color in the heatmap. The regions of similar color indicate that the element representations are either highly correlated or highly anticorrelated with each other. Furthermore, these regions are very similar between the two CBFVs. This is expected, since the CBFV features are based on physical properties of the elements. Thus, elements with similar physical properties will be more correlated while dissimilar elements will be more anticorrelated. Accordingly, the large colored regions typically correspond to similarities and dissimilarities between elements from families in the periodic table, such as alkali metals, alkaline earth metals, transition metals, metalloids and reactive nonmetals.

On the other hand, the element vectors from a DL model such as mat2vec do not exhibit such prominent behavior. Overall, the elements show less correlation with each other, and—with the exception of a few areas (to be discussed in later sections)—do not show large continuous regions of similar color. This is due to the fact that the starting element representations in DL models are randomly initialized and are not based on physical properties of the elements. These vector representation of the elements are only



**Fig. 1** Heatmaps of Pearson correlation matrices between element vectors featured using (a) Oliynyk, (b) Magpie, and (c) mat2vec element property feature sets. The  $x$  and  $y$  axis are labeled with the atomic numbers. Each cell at coordinate  $(x, y)$  represents the correlation between the corresponding elements with atomic numbers  $x$  and  $y$ . Blue represents a high correlation and red represents a high anticorrelation. For the interest of comparison, the heatmaps are truncated to the dimensions of the Oliynyk heatmap. Empty rows indicate that no element vector is available.

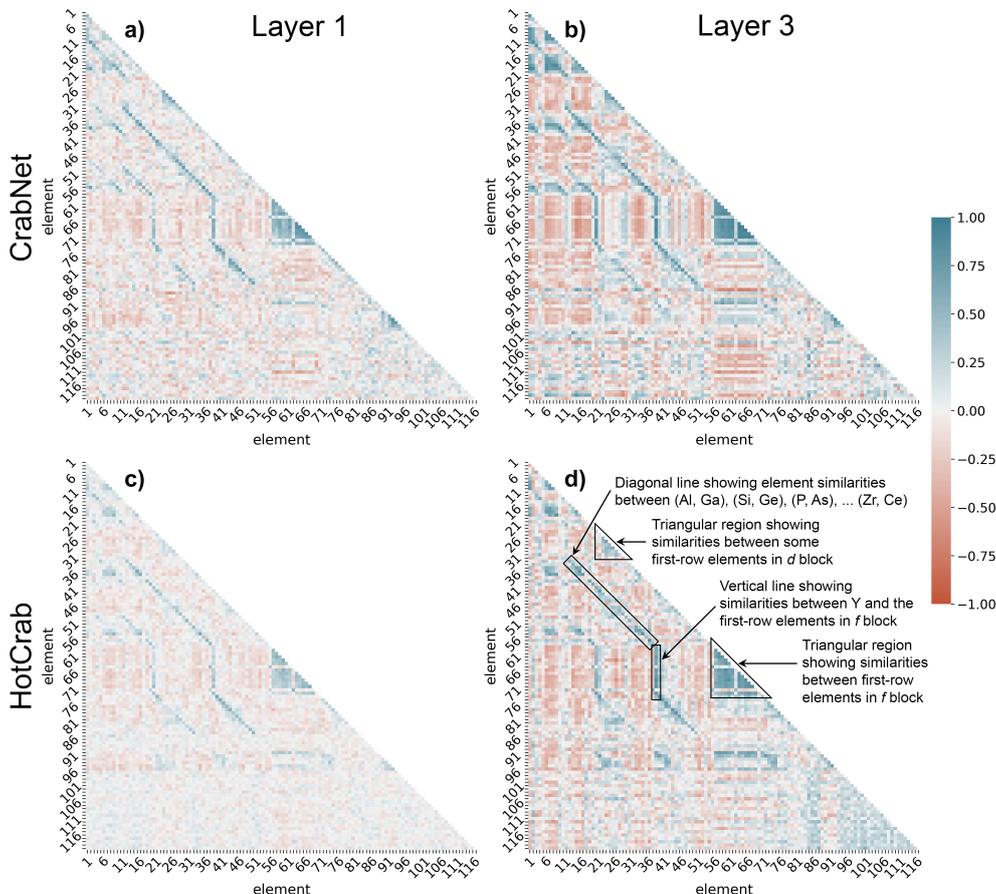
updated by the model throughout the training process using the training data. Thus, the correlation patterns that can be observed in this figure represent distinct patterns that the DL model has learned solely from the provided data.

We also note that a different number of element vectors are recorded in the feature sets. For the Oliynyk and Magpie CBFVs, only the elements up to uranium and berkelium are reported, respectively, while vectors up to the element oganesson are provided by mat2vec (please refer to Figure S-1 in the SI for the uncropped heatmaps). Particularly for the Oliynyk CBFV, some element vectors are missing, as visible by the empty rows in the heatmap. This disparity in the availability of element vectors between different CBFVs can be caused by reasons such as the instability or rarity of elements, lack of adequate information about the elements, or the inability to measure properties about the elements. The lack of element vectors in some material property feature sets can limit their applicability for certain tasks (such as when studying rare elements) and will be discussed in more detail in later sections.

In addition to learning element representations for a general purpose in materials science, such in the case of mat2vec, DL methods can also learn to relate element characteristics on a material property-specific basis. For example, element embeddings were extracted from the CrabNet and HotCrab models which were reproduced using the supplied model weights and the source code [57, 58]. The CrabNet and HotCrab models use mat2vec and onehot-encoded element features as the starting element representations, respectively. These features are then fine-tuned by the models for each of the 28 reported datasets. We extract one set of element embeddings from each layer of the models. Then, the Pearson correlation between the element vectors are plotted and shown in Figure 2.

In this work, we use the OQMD\_Bandgap dataset to demonstrate our findings. Additional example plots for other properties can be found in the SI. The OQMD datasets are widely used by researchers to evaluate model performance. For detailed information about the OQMD\_Bandgap dataset as well as information and discussion about the calculated values, please see the literature [59–61].

Here we can observe that both CrabNet and HotCrab are able to learn embeddings for each element of the periodic table, and that the correlations between the elements have a similar pattern, irrespective of the starting element representation (mat2vec or onehot). The observed correlation patterns are also similar to the mat2vec patterns as seen in Figure 1c. The ability of both CrabNet and HotCrab models to learn similar element embeddings despite having drastically different starting representations is encouraging, and further suggests that domain knowledge is not necessarily required for



**Fig. 2** Heatmaps of Pearson correlation matrices between element vectors extracted from CrabNet and HotCrab. These element representations are learned entirely from data. The  $x$  and  $y$  axis are labeled with the atomic numbers. Each cell at coordinate  $(x, y)$  represents the correlation between the corresponding elements with atomic numbers  $x$  and  $y$ . The top row (a and b) shows the correlations between embeddings from CrabNet and the bottom row (c and d) from HotCrab. The left and right columns represent the embeddings extracted from the first and last layer of the models, respectively. Blue represents a high correlation and red represents a high anticorrelation. In (d), some regions of interest are annotated.

element featurization if a sufficient quantity and quality of training data is available [18]. This finding is corroborated by the similarly good performance of both models across a wide range of material properties [28]. Interestingly, for deeper layers of the models (Figures 2b and d), more intense correlation patterns between the elements emerge. This is likely attributed to the self-attention based learning mechanism of the underlying CrabNet models. At each successive layer within the model, information about additional element-element interactions within the compound (*i.e.*, the chemical environment) are successively taken into account when updating the identity of an element within that compound. As a result, the deeper the layer within the model, the more complex the element interactions—and the element representations—become.

It is also interesting to note the diagonal and horizontal patterns which can be observed in all of the correlation matrices. For example, in Figure 2d there is a 45-degree diagonal, blue line that can be seen in the correlation matrix starting at the coordinates (13, 31) (corresponding to the element pair (Al, Ga)) and continuing until (40, 58) (corresponding to (Zr, Ce)). This line highlights the well-known periodic law

which states that elements with similar chemical properties fall into recurring periodic groups. Please refer to Figure S-2 for the enlarged version of the annotated heatmap and for correlation plots for other material properties. Another observation is the triangular region of high correlation between (57, 57) and (71, 71), which indicates that the first-row elements of the  $f$ -block are highly similar to each other. A similar triangular region can be observed between (23, 23) and (29, 29), indicating similarities between some first-row elements of the  $d$ -block. Lastly, the vertical blue line starting at the coordinates (39, 57) and continuing to (39, 71) indicate the chemical similarities between yttrium and the first-row elements of the  $f$ -block. These and other patterns can also be observed in the Oliylyk and Magpie CBFVs in Figure 1 as well. The ability of the CrabNet and HotCrab models to learn such chemical relationships which are comparable to hand-curated CBFVs based solely on the chemical formulae is exciting, and further reaffirms the finding that hand-engineering of features is not needed when training on big data [18].

Moreover, in Figure 2c we observe a distinct “border” at the element plutonium (with atomic number 94), where the correlation coefficients between the elements suddenly decrease and the patterns become less pronounced. Additional analysis of the OQMD\_Bandgap dataset showed that it does not contain any compounds with elements past plutonium. Due to the fact that the element representations are learned purely by the model from the dataset, their quality depends heavily on the quality of the dataset. Since the model performance depends on the quality of the element representations, by extension, it also then depends on the dataset quality [32].

We define element prevalence as the number of times a certain element has appeared as part of the compounds in a given dataset. When examining the OQMD\_Bandgap dataset, we note that there is an imbalance in element prevalence, with oxygen and copper appearing almost 1.5 times to twice as often, and fluorine, chlorine, bromine and iodine appearing only less than 0.1 times as often as the majority of the other elements in the dataset, respectively. This imbalance in element prevalence is even stronger for other datasets such as the aflow\_Egap, castelli, CritExam, mp\_e\_form and phonons datasets (see Figure S-3 in the SI for some example element prevalence plots).

### Quantifying dataset imbalance

The degree to which a dataset is imbalanced (otherwise referred to as its “evenness”) can be measured using the Shannon equitability index, which is a function of the Shannon entropy of the dataset [62–64]. Shannon entropy is widely used in information theory and can be used to characterize the degree of imbalance in a dataset [65,66]. The Shannon entropy  $H$  is defined in Equation 2 as:

$$H(X) = - \sum_{i=1}^k P(x_i) \log P(x_i) \quad (2)$$

where  $X$  is the set of discrete variables  $x_i \in \{x_1, \dots, x_n\}$ ,  $i$  is the class,  $P(x_i)$  is the proportional abundance of  $x_i$  and  $k$  is the total number of classes in the dataset.

For a dataset  $\mathcal{D}$  of  $n$  data occurrences and  $k$  distinct chemical elements (classes), each with counts  $c_i$ ,  $P(x_i) = \frac{c_i}{n}$  and the Shannon entropy can thus also be written as Equation 3:

$$H(\mathcal{D}) = - \sum_{i=1}^k \frac{c_i}{n} \log \left( \frac{c_i}{n} \right) \quad (3)$$

For continuity, we note that when  $c_i = 0$ , it means that no data sample is related to class  $i$  in the dataset, and therefore the multiplicand within the summation is defined

to be 0. Mathematically,  $\lim_{p \rightarrow 0^+} p \log(p) = 0$ . The maximum value of  $H(\mathcal{D})$  is  $\log(k)$ . This value occurs when all element classes in the dataset are observed at the same frequency (*i.e.*, the dataset is completely balanced). Therefore, the Shannon entropy  $H(\mathcal{D})$  is scaled by  $\log(k)$  to finally obtain the Shannon equitability index  $E(\mathcal{D})$ , which is defined in Equation 4 as:

$$E(\mathcal{D}) = \frac{H(\mathcal{D})}{\log(k)} \quad (4)$$

$E(\mathcal{D})$  ranges between 0 for a maximally imbalanced dataset and 1 for a maximally balanced dataset. The Shannon equitability indices are calculated for the 28 datasets examined in this work and are presented in Table 1. A plot showing the same information can be found in the SI (Figure S-4). For more information about the datasets, please refer to the CrabNet publication [28].

**Table 1** Shannon equitability indices calculated from the training data splits of the 28 reported datasets. Datasets were taken from [28].

material property dataset	equitability	material property dataset	equitability
castelli	0.823	aflow__ael_bulk_modulus_vrh	0.948
dielectric	0.864	aflow__ael_debye_temperature	0.948
elasticity_log10(G_VRH)	0.953	aflow__ael_shear_modulus_vrh	0.948
elasticity_log10(K_VRH)	0.953	aflow__agl_thermal_conductivity_300K	0.940
expt_gap	0.931	aflow__agl_thermal_expansion_300K	0.944
expt_is_metal	0.930	aflow__Egap	0.920
glass	0.771	aflow__energy_atom	0.917
jdft2d	0.872	CritExam__Ed	0.914
mp_e_form	0.913	CritExam__Ef	0.914
mp_gap	0.916	mp_bulk_modulus	0.923
mp_is_metal	0.916	mp_elastic_anisotropy	0.921
phonons	0.909	mp_e_hull	0.897
steels_yield	0.959	mp_mu_b	0.897
		mp_shear_modulus	0.921
		OQMD_Bandgap	0.976
		OQMD_Energy_per_atom	0.976
		OQMD_Formation_Enthalpy	0.976
		OQMD_Volume_per_atom	0.976

As can be seen in the table, the datasets studied in this work are not equally balanced in terms of element diversity. The more imbalanced a dataset is in terms of the element prevalence in the chemical compounds, the less likely the models will be able to adequately learn about the elements and their environments. The element embeddings learned for the infrequent elements will therefore be weaker and will not be able to capture as much information about these elements as compared to more frequently occurring elements. This leads to the observed weak correlation patterns between the less frequently seen elements beyond a certain cutoff atomic number in the datasets, as discussed earlier for Figure 2.

If the weakly learned elements are then encountered during inference time, the model will not be able to make an adequate prediction using the elements’ representations. Additionally, if certain elements or element combinations appear more frequently (majority classes) in the datasets as compared to other elements or combinations (minority classes), the model may be biased to better capture the behavior of majority classes at the expense of sacrificing performance on the minority classes. Such a dataset bias may appear in computational or experimental datasets due to the fact that some elements

are more commonly studied for certain material applications. On the other hand, certain elements (*e.g.*, rare or unstable elements) naturally occur less frequently and therefore are also contained in fewer compounds and datasets. Certain elements such as noble gases also rarely form compounds with other elements and are therefore rarely reported in materials datasets.

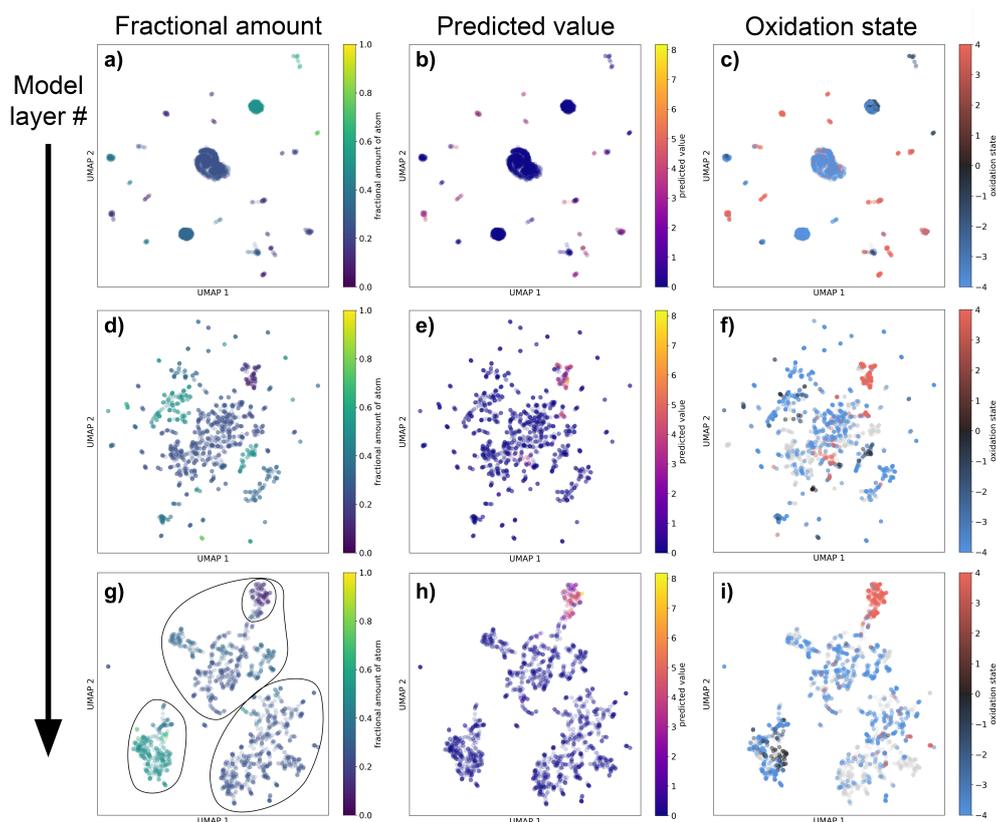
It is therefore important to implement data processing and modeling techniques to address biases as a result of dataset imbalance. Some example techniques include dataset re-sampling, generating synthetic data for imbalanced classes, implementing weighted loss functions that penalize errors for minority classes more, or using alternative loss functions and metrics to evaluate model performance [67,64,68]. Additionally, the model architecture can also be tailored to address dataset bias, and certain types of models (such as those based on self-attention or guided attention architectures) have an increased robustness against dataset bias [69,70].

Lastly, it is worthy to note that while most DL models learn element representations from structured materials datasets, methods such as word2vec and mat2vec use text mining and other natural language processing (NLP) techniques to learn the element embeddings from academic publications [71,20,72]. The data present in publications covers a much longer time period and contains a higher diversity in terms of types of compounds, material properties and applications studied. These data are in unstructured form and therefore cannot be used as training data for DL methods such as CrabNet; however, they can easily be used for word2vec and mat2vec. Therefore, text mining methods such as word2vec and mat2vec are able to learn from a much larger corpus of materials data and are not restricted by the availability of structured datasets. Accordingly, DL models such as CrabNet can benefit from the pre-trained element embeddings of mat2vec by fine-tuning the mat2vec embeddings to new tasks, thereby minimizing the impact of missing elements in the training dataset.

### Capturing the influence of chemical environments on element representations

In addition to learning the representations of each element, CrabNet and HotCrab can also capture the behavior of the elements when they are present in different chemical environments. Figure 3 shows the two-dimensional projections of the element vectors corresponding to the silicon atom from 2374 different silicon-containing compounds within the OQMD\_Bandgap test dataset. The silicon vectors are extracted from the transformed EDM tensors from HotCrab (a onehot-featurized version of CrabNet) and show the transformation of the silicon representations after they are passed through the three successive self-attention layers. For visualization, the vectors are projected down to two dimensions using the uniform manifold approximation and projection (UMAP) method [73]. The resulting points are plotted and colored by three parameters: (1) the fractional abundance of the element silicon in the compound, (2) the predicted property value of the compound (in this case, band gap), and (3) the oxidation state of silicon as predicted by Pymatgen [74]. For more information, please see the Methods.

As can be seen in the plots from the first layer (first row), there is a large number of distinct point clusters, with one major cluster near the center, two medium clusters above and below the center cluster, and many smaller clusters consisting of a few points. The larger clusters are formed because the initial representations of the silicon atoms are very similar to another (due to the learned element embedding of silicon). The similar silicon vectors are thus projected through UMAP into coordinates that lie close together, even though the silicon atoms are present in different chemical environments. We can observe as well that the clustering in layer one is mostly attributable to the fractional amount, since each cluster consists primarily of silicon points with the same fractional



**Fig. 3** Vector representations of the silicon element in 2374 different chemical environments and at different layers of the HotCrab model. Each point shows the model-internal representation of the silicon atom, after the information regarding the other atoms in the chemical environment have been introduced via HotCrab through the three attention layers (top row to bottom row). The points are colored by: (left column) the fractional abundance of silicon, (center column) the predicted value of the compound, and (right column) the predicted oxidation state of silicon, where gray points indicate that the oxidation state was unable to be predicted. Four clusters are outlined in the bottom-left plot.

amount. After the second layer, we observe that the points start to become separated into different and recognizable clusters. The clusters are no longer identifiable entirely based on the fractional amount of silicon, and clusters based on the predicted band gap value of the compound and oxidation state of silicon start to emerge. By the end of the third and last layer, we can observe four clusters that are distinguishable by the fractional amount of silicon, the predicted band gap, and the oxidation state of silicon (the clusters are outlined in Figure 3, bottom left).

More specifically, we observe that the cluster at the bottom-right side of the plot consists mainly of silicon with a fractional amount of around 0.15 to 0.3 (with a few points reaching 0.5), whereas the cluster near the bottom-left contains almost exclusively of silicon with fractional amounts of 0.5 plus a few points above 0.5. The cluster near the top contains regions of silicon with fractional amounts between 0.3 to 0.4 near the left and right, and around 0.2 to 0.3 in the middle. Near the top of this cluster, a smaller cluster is highlighted which consists mainly of silicon instances with low abundance, between 0.2 and 0. Please note that interactive versions of these plots can be found in the SI together with another example visualization plotted for the element chromium (Figure S-5).

In the predicted value plot of the last layer, we observe that only the small cluster near the top contains the silicon element in compounds with a non-zero band gap.

Similarly, when examining the oxidation state plot, we note that while most clusters contain a mixture of silicon atoms in several oxidation states, the same cluster near the top consists almost exclusively of silicon atoms in the +4 oxidation state and very few atoms in other oxidation states. Interestingly, while some compounds with silicon in the +4 state are visible in other clusters, these compounds have a zero band gap. This suggests that additional interactions between the elements were captured by HotCrab which lead to these compounds being correctly clustered together with other compounds with zero band gap.

These element behavior plots suggest that for silicon-containing compounds in the OQMD\_Bandgap dataset, the fractional amount and the oxidation state of the silicon atoms are important factors that together determine the band gap of the compounds. By cross-referencing the three plots, we can identify trends between the fractional amount and oxidation state of silicon and relate this information to the predicted band gap of the compounds. On the other hand, the clustering also suggests that there are other interactions between the elements in a compound which are currently not highlighted by the selected properties in Figure 3. It is our expectation that by examining these interactions, additional insight about the modeling process and element representations can be gained. Moreover, the findings from examining internal representations of elements in this way may suggest additional studies to further improve the understanding about the underlying phenomena governing materials behaviors. Note that while these visualizations were generated using HotCrab, similar results can be obtained using the CrabNet model.

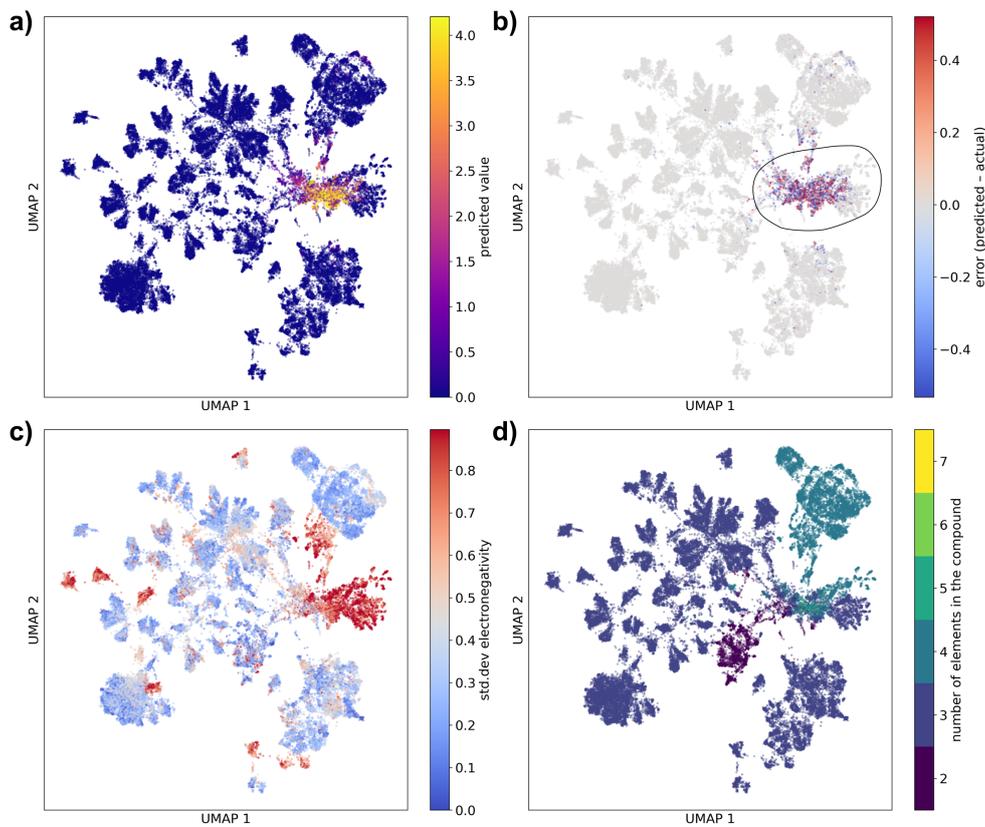
#### Capturing globally unique representations of chemical compounds

In addition to examining the behavior of individual elements in different chemical environments, we can also visualize all of the compounds in a given dataset to uncover additional insights. We extract the internal vector representation of all of the 51,242 compounds in the OQMD\_Bandgap test dataset from the last self-attention layer of HotCrab, perform dimensionality reduction using UMAP and finally visualize the compounds as shown in Figure 4. In addition to coloring the plots by the predicted value, prediction error, and number of distinct elements for the compounds, we also highlight the chemical trend between ionic to covalent bonding character within the compounds. This trend is revealed by calculating and visualizing the standard deviation of the Pauling electronegativities of the constituent atoms  $\sigma_\chi$  in a given compound [75] according to Equation 5:

$$\sigma_\chi = \sqrt{\frac{\sum_{i=1}^n (\chi_i - \bar{\chi})^2}{n - 1}} \quad (5)$$

where  $\chi_i$  is the Pauling electronegativity of each element  $i$  in the compound (totaling  $n$  elements), and  $\bar{\chi}$  is the average electronegativity of all elements in the compound. A higher  $\sigma_\chi$  signifies a more ionic bonding character, and a lower value signifies a more covalent bonding character.

Many clusters with varying sizes are visible in the figure. Some clusters are placed further apart, while some clusters are closer to, or are overlapping other clusters. In particular, the outlined cluster near the right of the figure is of particular interest. This is the only cluster where the compounds with a non-zero band gap are located, as is visible from Figure 4a. Additionally, it is also within this cluster that HotCrab makes the largest errors when predicting the band gap value, as seen in Figure 4b. For the other compounds, the prediction errors of HotCrab are close to zero. Even through a small proportion of model predictions have larger errors, the overall model performance is very



**Fig. 4** Global representations of the 51,242 compounds in the OQMD\_Bandgap test dataset, extracted from layer three of HotCrab, embedded down to two dimensions using UMAP and colored by the parameters: (a) the predicted value of the compound (band gap); (b) the prediction error ( $\hat{y} - y$ ); (c) the bond character of the compounds ranging from more covalent (blue) to more ionic (red) as measured by the standard deviations in the Pauling electronegativities of the constituent elements; and (d) the number of distinct elements in the compound. A cluster of interest is outlined in the plot at the top-right.

good and is comparable with, or better than, other state-of-the-art models [28]. This superior performance of CrabNet and HotCrab models when predicting properties with a defined cutoff (such as the cutoff of 0 eV in this case for band gap) is likely attributed to the prediction of element-logits in the modeling process. These element-logits are used to weight the final model predictions in CrabNet and HotCrab to improve the model accuracy [28].

Notably, we also observe from Figures 4a, c and d that the band gap only partially depends on the bond nature of the compound and on the number of unique elements in the compound. While most of the compounds in the cluster of interest exhibit more ionic bond characters, there are also other clusters with similar bond character that do not have a non-zero band gap. Similarly, it appears that the compounds with a non-zero band gap mainly contain four or five unique elements; however, there are also other compounds with these numbers of unique elements which have a zero band gap.

Here we do note that while UMAP can reveal structures and patterns within high-dimensional data, it generally emphasizes local structure at the expense of global structure. Therefore, for the UMAP visualizations shown in this work, it is more appropriate to interpret the local structure (*e.g.*, the elements or compounds present within individual clusters in Figures 3 and 4) than the global structure. While the number of local

neighbors considered can be specified as a hyperparameter in UMAP, a trade-off is made between preserving local versus global structure. Therefore, the distances between elements and compounds within a single cluster are more meaningful than inter-cluster distances in the UMAP visualizations. Lastly, we note that while these visualizations were generated based on the test dataset using HotCrab, similar results can be obtained using CrabNet or the training dataset.

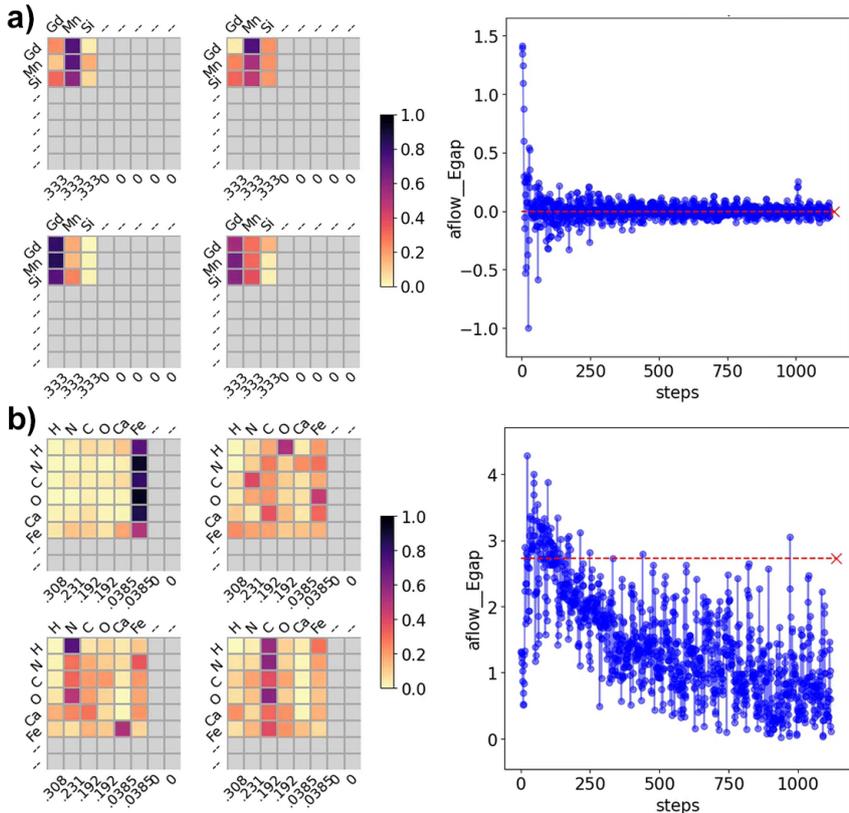
### Visualizing the training progress

Beyond visualizing the element and compound representations from CrabNet after training, it is also possible to access the self-attention matrices of the CrabNet encoding layers to observe the model learning process during training. The attention matrices (commonly referred to as the attention maps) contain information regarding how each element (rows) is influenced by all other elements in the compound as well as itself (columns). The values in the attention maps are the attention scores and are used in the encoder to update the element representations. An attention score of zero means that the element in the column is completely ignored when updating the element’s representation in that row. Conversely, a score of one means that the entire update is based solely on that column’s element.

In the CrabNet publication [28], example attention maps were shown for compounds after the model has finished training. Here, we extend this approach by visualizing the CrabNet attention maps during the model training process in the form of attention video clips (see SI files for examples). This is achieved by saving the attention matrices from the model encoder layers after every mini-step in the training process and generating a video to show the learning progress. Figure 5 shows a snapshot of two example attention videos obtained at the end of model training. The attention maps from the first encoding layer of CrabNet are plotted as heatmaps in the left column, while the right column shows the predicted values from the model versus the target value at every mini-step. This process is performed at every mini-step in the training process, and the resulting plots are merged into a video clip which shows the learning progress of the model throughout training.

From the attention maps, we can observe that some elements are considered less relevant in the determination of the material property, whereas some elements are considered very relevant. Also we can note that individual attention heads pay attention to different element-element interactions in the compound, as is visible by the significantly different attention patterns in the plots. Throughout the training process, the attention pattern for each head remains relatively fixed after a few mini-steps, indicating that the model discovers a pattern for recognizing inter-element interactions early on in the training process, which it then continues to refine as more training steps are taken.

For the top compound, we can observe that while the model initially over- and underestimates the property value early on in the training, it learns to correct the error and finally achieves a low prediction error towards the end of training. Conversely, for the bottom compound, we observe that while the model initially correctly estimates the property value of the compound, the predicted value decreases and the estimation error increases throughout training, with the error finally plateauing towards the end of the training. By examining the attention heatmaps for this compound, we notice that attention head 1 shows a significantly different behavior as compared to the other attention heads. It dedicates almost all of its attention to the element iron, while the other attention heads capture many more inter-element interactions. It may be interesting to investigate further to find out if CrabNet is misrepresenting the interactions from the iron element with the other elements and thus making the prediction error, or if another phenomenon is contributing to the prediction error on this compound.



**Fig. 5** Snapshots of attention videos for observing the training progress of CrabNet using two example compounds (a)  $\text{Gd}_1\text{Mn}_1\text{Si}_1$  and (b)  $\text{C}_5\text{Ca}_1\text{Fe}_1\text{H}_8\text{N}_6\text{O}_5$  from the validation data split of the `aflow_Egap` dataset. The left plots show the attention maps of the four attention heads at the first attention layer, where the  $x$  axis of each heatmap is labeled with the fractional amount of the elements and the other axes are labeled with the element symbol. The right plots show the model predictions (blue) for the compounds, evaluated after each training mini-step throughout the whole training process. The true property value (target) is represented with the red “X” and the dotted line.

By observing the element groups and inter-elemental interactions that CrabNet pays attention to for each material property throughout the training process, we may be able to gain additional insight about which relevant elements and interactions contribute significantly to the material property. Similarly, in the case where the model does not make a good property prediction or fails to learn a specific material property, these attention videos can be informative in showing when, where, and how the model fails. Additionally, since the element representations in a compound are updated according to the attention scores, it would be interesting to train CrabNet on material properties where the property has a high sensitivity to changes in elemental prevalence. An example of this is in the case of dopants, where a small change in the dopant amount can significantly influence a material’s electrical [15, 76, 77], mechanical [78–80, 17], and thermal properties [81–84]. Finally, it may be interesting to expand the studied materials to include co-doped materials and use the attention videos to visualize the complex inter-elemental interactions between the co-dopants and the host elements.

## Conclusion

In this work, we examined the CrabNet model through the use of several built-in model interpretability methods in order to visualize the data featurization and modeling process. We demonstrated that CrabNet can adequately capture the chemical behavior of compounds in a dataset by using the vector representations of their constituent elements. The element representations can be learned entirely from the training data on a per-property basis, and contain rich information about the elements and their chemical trends. Additionally, we examined dataset imbalance, its relation to the quality of learned representations, and the limitations that imbalanced datasets may ultimately impose on the modeling processes.

The element and compound vectors can be projected using UMAP into distinguishable clusters which can then be visualized and characterized by the element stoichiometry, local chemical environment and oxidation state of the elements, or by the bond behavior of the compounds. Lastly, the examination of the self-attention matrices during model training through the use of attention videos can be used to further understand the modeling process, debug potential model or dataset errors, or gain additional insights about chemical interactions within a given compound.

The model interpretability techniques presented in this work will enable materials science practitioners to not only visualize a specific element’s behavior within different chemical environments, but also to obtain a global view of the chemical compounds, behaviors and trends within a larger dataset. The ability of CrabNet to adequately model and express the complex chemical behaviors and interactions of elements and compounds based solely on learning from data is encouraging. With the addition of model interpretability methods to CrabNet, the findings and intuitions presented in this work may lead to further insightful and interesting research. Specifically, we believe that follow-up works may fall into one of these three general directions:

1. **Learning and representing elements and compounds.** Our work has shown that it is possible to visualize CrabNet’s internal representations of elements and compounds via techniques such as UMAP. However, it would be interesting to further investigate why CrabNet’s representations of some of these elements or compounds lead to them being placed into the same cluster or not, despite the fact that these elements and compounds are similar to each other in terms of identity and/or chemical environment. This may also be combined with a more detailed examination of the attention videos and how the attention mechanism in CrabNet leads to the updating of the element representations for each compound.
2. **Examination of individual attention head behaviors.** This work used the EDM (element-derived matrix) data from CrabNet to examine the element and compound representations within CrabNet. CrabNet utilizes four self-attention heads to model element-element interactions, the results of which are then concatenated and transformed back to an updated EDM matrix. As such, the EDM is a pooled representation of the compounds. It would be interesting to further examine the per-head modeling of the compounds, as it has been shown that each head can capture different types of inter-element interactions and thus may give additional insight to the modeling process within CrabNet.
3. **Discovery of additional inter-element interactions.** From the analyses presented in this study, it is clear that while some changes in the material property (*e.g.*, band gap) can be explained by certain properties of the compounds (such as element stoichiometry, number of unique elements, and/or bond character), there are additional behaviors that govern the material property. These additional interactions are also adequately modeled by CrabNet, since it can predict a wide range of material properties with low errors. Examining the modeling process of these behaviors

within CrabNet may lead to an improved understanding of the complex phenomena underlying material properties.

Further research to answer these and subsequent questions may allow us to gain additional insights about the behaviors and properties of elements and materials, improve our understanding of models such as CrabNet, increase our confidence in the use of data-driven methods, and ultimately, accelerate the adoption of deep learning and machine learning in materials science.

## Methods

### Adaptation of CrabNet model

The CrabNet model and material property datasets as originally reported were used as the basis for this study [28]. Fully trained model weights for both CrabNet and HotCrab were obtained from [57]. In order to obtain the EDMs containing the elements and compounds data used in this study, custom function hooks were implemented in PyTorch. These hooks were attached to the CrabNet model architecture to allow access to the model-internal data during training and inference.

The source code as well as the data that were used and generated in this study can be found on the updated CrabNet GitHub repository [58] (the GitHub will be updated within a few days of manuscript acceptance). In addition, we provide detailed instructions for the use and reproduction of our reported results. Please note that due to the prohibitively large size of the stored attention matrices used in the attention videos, it is not possible to provide these for download. However, instructions and scripts are provided for generating these matrices and videos.

All experiments, unless otherwise noted, were performed on a workstation equipped with an Intel i7-8700K CPU, 32 GB of DDR4 RAM, and one Nvidia RTX 2080 GPU.

### Element Embeddings

Element embeddings for pure elements were generated on a per-property basis. To do this, an EDM consisting of all of the elements from hydrogen to oganesson was generated (with each row representing one element). Then, for each material property, the corresponding CrabNet or HotCrab model was loaded and the model hooks attached. The EDM was then passed through the network and the modified EDM at the output of the element embedding layer was obtained and detached from the model graph. This resulting EDM contains the property-specific element embeddings of all of the elements. Thus, each element was represented by a vector with the shape  $(1, d_{\text{model}})$ , where  $d_{\text{model}}$  is the size of the embedding. Element embeddings for Oliynyk, Magpie, and mat2vec were obtained from the original publications [18].

### Compound Embeddings

Compound embeddings were obtained in a similar fashion to element embeddings. Instead of generating an EDM from pure elements, the EDMs were generated from the actual chemical formulae from the datasets and collated in batches using the model data loader. Model hooks were then attached to the CrabNet and HotCrab models and enabled during model inference. The transformed EDMs after each of the three self-attention layers of the CrabNet models were then collected.

The obtained compound EDMs have the shape of  $(n_{\text{compounds}}, n_{\text{elements}}, d_{\text{model}})$ , where  $n_{\text{compounds}}$  is the total number of compounds in the dataset,  $n_{\text{elements}}$  is the maximum number of elements per compound, and  $d_{\text{model}}$  is the size of the embedding. Thus, each compound in the EDM is represented by one tensor slice with the dimensions  $(1, n_{\text{elements}}, d_{\text{model}})$ . Due to the fact that different compounds within the same dataset may contain a different number of elements, the extra rows of the EDMs were zero-filled to indicate no elements present. In order to ensure that the compound embeddings are comparable with each other using UMAP, the three-dimensional compound EDMs were collapsed to two dimensions  $(n_{\text{compounds}}, 1, d_{\text{model}})$  by calculating summary statistics (such as sum, range, variance) of the EDM columns across the elements dimension.

### Dimensionality Reduction

CrabNet uses vectors with a  $d_{\text{model}}$  dimension of 512 to represent chemical elements and compounds in the input data. It would be infeasible to try to visualize all 512 dimensions. Therefore, dimensionality reduction was applied to the vector representations to transform the vectors into two-dimensional space for visualization.

Three common methods for dimensionality reduction were tested: principal component analysis (PCA), t-distributed stochastic neighbor embedding (t-SNE), and uniform manifold approximation and projection (UMAP) [85, 86, 73]. Compared to t-SNE and PCA, UMAP revealed more visually distinct clusters for the data presented in this work. Therefore, UMAP was chosen as the dimensionality reduction method. The random seed was fixed so that each initialization of the UMAP method produces the same results. For element embeddings, the rows of the EDMs with dimensions  $(1, d_{\text{model}})$  are transformed using UMAP. For the compound embeddings, the matrices corresponding to each compound were first collapsed as described above, and the resulting representations with dimensions  $(1, d_{\text{model}})$  for each compound were transformed using UMAP.

### Oxidation State Estimation

Oxidation states for elements in the compounds were estimated using the Pymatgen package (version 2022.0.8) using the chemical formulae of the compounds. The built-in functions for assigning oxidation states were used, which are based on charge-balancing heuristics and use the most probable oxidation states as determined based on the compounds in the Inorganic Crystal Structure Database [74].

### Attention Video Generation

Custom function hooks were programmed and attached to a newly-initialized CrabNet model. During training of CrabNet, the attention matrices of every CrabNet encoder layer was extracted from the model and saved into a compressed Zarr array on disk. The model predictions for the properties were also generated and saved. This procedure is performed after every mini-step during the training process (corresponding to each mini-batch of data). The plots were then generated for each mini-step and merged together using the software FFMPEG to create the attention videos. Due to the large amount of storage and computing power required to store and process the attention matrices, these tasks were performed on a high-performance computing cluster.

## Acknowledgments

The authors thank the Berlin International Graduate School in Model and Simulation based Research as well as the German Academic Exchange Service RISE program for their financial support. Special thanks is given to Dr. Steven K. Kauwe, Pay Gießelmann and Joris Weigert for the insightful discussions.

Computing resources were graciously provided by the HPC-Cluster at the Institut für Mathematik, Technische Universität Berlin, by the HPC Resource in the Core Facility for Advanced Research Computing at the Case Western Reserve University as well as by the Google TPU Research Cloud (TRC) program.

In addition, the authors express their gratitude to the open-source software community for developing the excellent tools used in this research, including but not limited to Python, Pandas, NumPy, matplotlib, scikit-learn, PyTorch, Zarr, and FFMPEG.

## Conflict of interest

On behalf of all authors, the corresponding author states that there is no conflict of interest.

## References

1. R. Ramprasad, R. Batra, G. Pilania, A. Mannodi-Kanakkithodi, and C. Kim, "Machine learning in materials informatics: Recent applications and prospects," *npj Computational Materials*, vol. 3, no. 1, p. 60, 2017.
2. J. Schmidt, M. R. G. Marques, S. Botti, and M. A. L. Marques, "Recent advances and applications of machine learning in solid-state materials science," *npj Computational Materials*, vol. 5, no. 1, p. 83, 2019.
3. C. P. Gomes, B. Selman, and J. M. Gregoire, "Artificial intelligence for materials discovery," *MRS Bulletin*, vol. 44, no. 7, pp. 538–544, 2019.
4. O. Isayev, A. Tropsha, and S. Curtarolo, eds., *Materials Informatics: Methods, Tools, and Applications*. Wiley, 2019.
5. B. DeCost, J. R. Hattrick-Simpers, Z. Trautt, A. G. Kusne, E. Campo, and M. L. Green, "Scientific AI in materials science: a path to a sustainable and scalable paradigm," *Machine Learning: Science and Technology*, 2020.
6. H. S. Stein and J. M. Gregoire, "Progress and prospects for accelerating materials science with automated and autonomous workflows," *Chemical Science*, vol. 10, no. 42, pp. 9640–9649, 2019.
7. D. Morgan and R. Jacobs, "Opportunities and Challenges for Machine Learning in Materials Science," *Annual Review of Materials Research*, vol. 50, no. 1, pp. 71–103, 2020.
8. C. L. Zitnick, L. Chanussot, A. Das, S. Goyal, J. Heras-Domingo, C. Ho, W. Hu, T. Lavril, A. Palizhati, M. Riviere, M. Shuaibi, A. Sriram, K. Tran, B. Wood, J. Yoon, D. Parikh, and Z. Ulissi, "An Introduction to Electrocatalyst Design using Machine Learning for Renewable Energy Storage," 2020-10-14.
9. T. D. Sparks, S. K. Kauwe, M. E. Parry, A. Mansouri Tehrani, and J. Brgoch, "Machine Learning for Structural Materials," *Annual Review of Materials Research*, vol. 50, no. 1, 2020.
10. G. Pilania, "Machine learning in materials science: From explainable predictions to autonomous design," *Computational Materials Science*, vol. 193, p. 110360, 2021.
11. A. O. Oliynyk, E. Antono, T. D. Sparks, L. Ghadbeigi, M. W. Gaultois, B. Meredig, and A. Mar, "High-Throughput Machine-Learning-Driven Synthesis of Full-Heusler Compounds," *Chemistry of Materials*, vol. 28, no. 20, pp. 7324–7331, 2016.
12. L. Ward, A. Agrawal, A. Choudhary, and C. Wolverton, "A general-purpose machine learning framework for predicting properties of inorganic materials," *npj Computational Materials*, vol. 2, no. 1, p. 16028, 2016.
13. G. Pilania, A. Mannodi-Kanakkithodi, B. P. Uberuaga, R. Ramprasad, J. E. Gubernatis, and T. Lookman, "Machine learning bandgaps of double perovskites," *Scientific Reports*, vol. 6, p. 19375, 2016.
14. A. Dunn, Q. Wang, A. Ganose, D. Dopp, and A. Jain, "Benchmarking materials property prediction methods: the Matbench test set and Automatminer reference algorithm," *npj Computational Materials*, vol. 6, no. 1, p. 138, 2020.

15. S. K. Kauwe, J. Graser, R. J. Murdock, and T. D. Sparks, "Can machine learning find extraordinary materials?," *Computational Materials Science*, vol. 174, p. 109498, 2020.
16. J. Graser, S. K. Kauwe, and T. D. Sparks, "Machine Learning and Energy Minimization Approaches for Crystal Structure Predictions: A Review and New Horizons," *Chemistry of Materials*, vol. 30, no. 11, pp. 3601–3612, 2018.
17. A. Mansouri Tehrani, A. O. Oliynyk, M. Parry, Z. Rizvi, S. Couper, F. Lin, L. Miyagi, T. D. Sparks, and J. Brgoch, "Machine Learning Directed Search for Ultraincompressible, Superhard Materials," *Journal of the American Chemical Society*, vol. 140, no. 31, pp. 9844–9853, 2018.
18. R. J. Murdock, S. K. Kauwe, A. Y.-T. Wang, and T. D. Sparks, "Is Domain Knowledge Necessary for Machine Learning Materials Properties?," *Integrating Materials and Manufacturing Innovation*, vol. 9, no. 3, pp. 221–227, 2020.
19. K. Choudhary, B. DeCost, and F. Tavazza, "Machine learning with force-field-inspired descriptors for materials: Fast screening and mapping energy landscape," *Physical Review Materials*, vol. 2, no. 8, p. 083801, 2018.
20. V. Tshitoyan, J. Dagdelen, L. Weston, A. Dunn, Z. Rong, O. Kononova, K. A. Persson, G. Ceder, and A. Jain, "Unsupervised word embeddings capture latent knowledge from materials science literature," *Nature*, vol. 571, no. 7763, pp. 95–98, 2019.
21. S. K. Kauwe, J. Graser, A. Vazquez, and T. D. Sparks, "Machine Learning Prediction of Heat Capacity for Solid Inorganics," *Integrating Materials and Manufacturing Innovation*, vol. 7, no. 2, pp. 43–51, 2018.
22. S. K. Kauwe, T. Welker, and T. D. Sparks, "Extracting Knowledge from DFT: Experimental Band Gap Predictions Through Ensemble Learning," *Integrating Materials and Manufacturing Innovation*, vol. 9, no. 3, pp. 213–220, 2020.
23. D. Jha, L. Ward, A. Paul, W.-K. Liao, A. Choudhary, C. Wolverton, and A. Agrawal, "ElemNet: Deep Learning the Chemistry of Materials From Only Elemental Composition," *Scientific Reports*, vol. 8, no. 1, p. 17593, 2018.
24. T. Xie and J. C. Grossman, "Crystal Graph Convolutional Neural Networks for an Accurate and Interpretable Prediction of Material Properties," *Physical Review Letters*, vol. 120, no. 14, p. 145301, 2018.
25. C. Chen, W. Ye, Y. Zuo, C. Zheng, and S. P. Ong, "Graph Networks as a Universal Machine Learning Framework for Molecules and Crystals," *Chemistry of Materials*, vol. 31, no. 9, pp. 3564–3572, 2019.
26. J. Klicpera, S. Giri, J. T. Margraf, and S. Günnemann, "Fast and Uncertainty-Aware Directional Message Passing for Non-Equilibrium Molecules," 2020-11-28.
27. B. DeCost and K. Choudhary, "Atomistic Line Graph Neural Network for Improved Materials Property Predictions," 2021-06-03.
28. A. Y.-T. Wang, S. K. Kauwe, R. J. Murdock, and T. D. Sparks, "Compositionally restricted attention-based network for materials property predictions," *npj Computational Materials*, vol. 7, no. 1, p. 77, 2021.
29. R. E. A. Goodall and A. A. Lee, "Predicting materials properties without crystal structure: deep representation learning from stoichiometry," *Nature Communications*, vol. 11, no. 1, p. 6280, 2020.
30. S. Kong, D. Guevarra, C. P. Gomes, and J. M. Gregoire, "Materials representation and transfer learning for multi-property prediction," *Applied Physics Reviews*, vol. 8, no. 2, p. 021409, 2021.
31. C. L. Clement, S. K. Kauwe, and T. D. Sparks, "Benchmark AFLOW Data Sets for Machine Learning," *Integrating Materials and Manufacturing Innovation*, vol. 9, no. 2, pp. 153–156, 2020.
32. A. Y.-T. Wang, R. J. Murdock, S. K. Kauwe, A. O. Oliynyk, A. Gurlo, J. Brgoch, K. A. Persson, and T. D. Sparks, "Machine Learning for Materials Scientists: An Introductory Guide toward Best Practices," *Chemistry of Materials*, vol. 32, no. 12, pp. 4954–4965, 2020.
33. A. N. Henderson, S. K. Kauwe, and T. D. Sparks, "Benchmark datasets incorporating diverse tasks, sample sizes, material systems, and data heterogeneity for materials informatics," *Data in Brief*, vol. 37, p. 107262, 2021.
34. B. Meredig, "Industrial materials informatics: Analyzing large-scale data to solve applied problems in R&D, manufacturing, and supply chain," *Current Opinion in Solid State & Materials Science*, vol. 21, no. 3, pp. 159–166, 2017.
35. Z. C. Lipton, "The Mythos of Model Interpretability," *Queue*, vol. 16, no. 3, pp. 31–57, 2018.
36. L. Himanen, A. Geurts, A. S. Foster, and P. Rinke, "Data-Driven Materials Science: Status, Challenges, and Perspectives," *Advanced Science*, vol. 6, no. 21, p. 1900808, 2019.
37. C. Rudin, "Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead," *Nature Machine Intelligence*, vol. 1, no. 5, pp. 206–215, 2019.
38. I. Kolyshkina and S. Simoff, "Interpretability of Machine Learning Solutions in Industrial Decision Engineering," in *Data Mining* (T. D. Le, K.-L. Ong, Y. Zhao, W. H. Jin, S. Wong, L. Liu, and G. Williams, eds.), vol. 1127 of *Communications in Computer and Information Science*, pp. 156–170, Singapore: Springer Singapore, 2019.
39. P. Linardatos, V. Papastefanopoulos, and S. Kotsiantis, "Explainable AI: A Review of Machine Learning Interpretability Methods," *Entropy (Basel, Switzerland)*, vol. 23, no. 1, p. 18, 2021.

40. L. H. Gilpin, D. Bau, B. Z. Yuan, A. Bajwa, M. Specter, and L. Kagal, "Explaining Explanations: An Overview of Interpretability of Machine Learning," in *2018 IEEE 5th International Conference on Data Science and Advanced Analytics (DSAA)*, pp. 80–89, IEEE, 2018.
41. D. Smilkov, N. Thorat, C. Nicholson, E. Reif, F. B. Viégas, and M. Wattenberg, "Embedding Projector: Interactive Visualization and Interpretation of Embeddings," 2016-11-16.
42. S. Liu, P.-T. Bremer, J. J. Thiagarajan, V. Srikumar, B. Wang, Y. Livnat, and V. Pascucci, "Visual Exploration of Semantic Relationships in Neural Word Embeddings," *IEEE Transactions on Visualization and Computer Graphics*, vol. 24, no. 1, pp. 553–562, 2018.
43. B. van Aken, B. Winter, A. Löser, and F. A. Gers, "VisBERT: Hidden-State Visualizations for Transformers," in *Companion Proceedings of the Web Conference 2020* (A. E. F. Seghrouchni, G. Sukthankar, T.-Y. Liu, and M. van Steen, eds.), (New York, NY, USA), pp. 207–211, ACM, 2020.
44. O. Vinyals, I. Babuschkin, W. M. Czarnecki, M. Mathieu, A. Dudzik, J. Chung, D. H. Choi, R. Powell, T. Ewalds, P. Georgiev, J. Oh, D. Horgan, M. Kroiss, I. Danihelka, A. Huang, L. Sifre, T. Cai, J. P. Agapiou, M. Jaderberg, A. S. Vezhnevets, R. Leblond, T. Pohlen, V. Dalibard, D. Budden, Y. Sulsky, J. Molloy, T. L. Paine, C. Gulcehre, Z. Wang, T. Pfaff, Y. Wu, R. Ring, D. Yogatama, D. Wünsch, K. McKinney, O. Smith, T. Schaul, T. Lillicrap, K. Kavukcuoglu, D. Hassabis, C. Apps, and D. Silver, "Grandmaster level in StarCraft II using multi-agent reinforcement learning," *Nature*, vol. 575, no. 7782, pp. 350–354, 2019.
45. E. Puiutta and E. M. S. P. Veith, "Explainable Reinforcement Learning: A Survey," in *Machine Learning and Knowledge Extraction* (A. Holzinger, P. Kieseberg, A. M. Tjoa, and E. Weippl, eds.), vol. 12279 of *Lecture Notes in Computer Science*, pp. 77–95, Cham: Springer International Publishing, 2020.
46. A. Heuillet, F. Couthouis, and N. Díaz-Rodríguez, "Explainability in deep reinforcement learning," *Knowledge-Based Systems*, vol. 214, p. 106685, 2021.
47. S. Lopuschkin, *Opening the machine learning black box with Layer-wise Relevance Propagation*. PhD thesis, Technische Universität Berlin, Berlin, Germany, 2018.
48. H. Chefer, S. Gur, and L. Wolf, "Transformer Interpretability Beyond Attention Visualization," 2020-12-17.
49. J. Chen, Y. Lu, Q. Yu, X. Luo, E. Adeli, Y. Wang, Le Lu, A. L. Yuille, and Y. Zhou, "TransUNet: Transformers Make Strong Encoders for Medical Image Segmentation," 2021-02-08.
50. S. Khan, M. Naseer, M. Hayat, S. W. Zamir, F. S. Khan, and M. Shah, "Transformers in Vision: A Survey," 2021-01-04.
51. B. Kailkhura, B. Gallagher, S. Kim, A. Hiszpanski, and T. Y.-J. Han, "Reliable and explainable machine-learning methods for accelerated material discovery," *npj Computational Materials*, vol. 5, no. 1, p. 221, 2019.
52. R. Roscher, B. Bohn, M. F. Duarte, and J. Garcke, "Explainable Machine Learning for Scientific Insights and Discoveries," *IEEE Access*, vol. 8, pp. 42200–42216, 2020.
53. M. T. Ribeiro, S. Singh, and C. Guestrin, "Why Should I Trust You?: Explaining the Predictions of Any Classifier," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining – KDD '16* (B. Krishnapuram, M. Shah, A. Smola, C. Aggarwal, D. Shen, and R. Rastogi, eds.), (New York, NY, USA), pp. 1135–1144, ACM Press, 2016.
54. S. Lundberg and S.-I. Lee, "A Unified Approach to Interpreting Model Predictions," 2017.
55. L. S. Shapley, "A Value for n-Person Games," in *Contributions to the Theory of Games (AM-28), Volume II* (H. W. Kuhn and A. W. Tucker, eds.), Annals of Mathematics Studies, pp. 307–318, Princeton, NJ: Princeton University Press, 1953.
56. T. Hastie, R. Tibshirani, and J. H. Friedman, *The elements of statistical learning: Data mining, inference, and prediction*. Springer Series in Statistics, New York, NY: Springer, 2nd ed. ed., 2009.
57. A. Y.-T. Wang, S. K. Kauwe, R. J. Murdock, and T. D. Sparks, "Trained network weights for the paper "Compositionally restricted attention-based network for materials property predictions (CrabNet)," 2021.
58. A. Y.-T. Wang and S. K. Kauwe, "Online GitHub repository for the paper "Compositionally-Restricted Attention-Based Network for Materials Property Prediction," 2020.
59. J. E. Saal, S. Kirklin, M. Aykol, B. Meredig, and C. Wolverton, "Materials Design and Discovery with High-Throughput Density Functional Theory: The Open Quantum Materials Database (OQMD)," *JOM*, vol. 65, no. 11, pp. 1501–1509, 2013.
60. S. Kirklin, J. E. Saal, B. Meredig, A. Thompson, J. W. Doak, M. Aykol, S. Rühl, and C. Wolverton, "The Open Quantum Materials Database (OQMD): assessing the accuracy of DFT formation energies," *npj Computational Materials*, vol. 1, no. 1, p. 15010, 2015.
61. V. I. Hegde, C. K. H. Borg, Z. d. Rosario, Y. Kim, M. Hutchinson, E. Antono, J. Ling, P. Saxe, J. E. Saal, and B. Meredig, "Reproducibility in high-throughput density functional theory: a comparison of AFLOW, Materials Project, and OQMD," 2020-07-04.
62. J. A. Bonachela, H. Hinrichsen, and M. A. Muñoz, "Entropy estimates of small data sets," *Journal of Physics A: Mathematical and Theoretical*, vol. 41, no. 20, p. 202001, 2008.
63. C. Hong, R. Ghosh, and S. Srinivasan, "Dealing with Class Imbalance using Thresholding," 2016-07-10.

64. M. A. U. H. Tahir, S. Asghar, A. Manzoor, and M. A. Noor, "A Classification Model For Class Imbalance Dataset Using Genetic Programming," *IEEE Access*, vol. 7, pp. 71013–71037, 2019.
65. C. E. Shannon, "A Mathematical Theory of Communication," *Bell System Technical Journal*, vol. 27, no. 3, pp. 379–423, 1948.
66. C. E. Shannon, "A Mathematical Theory of Communication," *Bell System Technical Journal*, vol. 27, no. 4, pp. 623–656, 1948.
67. Y. Li and N. Vasconcelos, "REPAIR: Removing Representation Bias by Dataset Resampling," in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (CVPR Editors, ed.), pp. 9564–9573, IEEE, 2019.
68. C. Esposito, G. A. Landrum, N. Schneider, N. Stiefl, and S. Riniker, "GHOST: Adjusting the Decision Threshold to Handle Imbalanced Data in Machine Learning," *Journal of Chemical Information and Modeling*, vol. 61, no. 6, pp. 2623–2640, 2021.
69. K. Li, Z. Wu, K.-C. Peng, J. Ernst, and Y. Fu, "Tell Me Where to Look: Guided Attention Inference Network," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9215–9223, IEEE, 2018.
70. A. C. Rodriguez, S. D'Aronco, K. Schindler, and J. D. Wegner, "Privileged Pooling: Better Sample Efficiency Through Supervised Attention," 2020-03-20.
71. E. Kim, K. Huang, A. Tomala, S. Matthews, E. Strubell, A. Saunders, A. McCallum, and E. Olivetti, "Machine-learned and codified synthesis parameters of oxide materials," *Scientific Data*, vol. 4, p. 170127, 2017.
72. L. Weston, V. Tshitoyan, J. Dagdelen, O. Kononova, A. Trewartha, K. A. Persson, G. Ceder, and A. Jain, "Named Entity Recognition and Normalization Applied to Large-Scale Information Extraction from the Materials Science Literature," *Journal of Chemical Information and Modeling*, 2019.
73. L. McInnes, J. Healy, N. Saul, and L. Großberger, "UMAP: Uniform Manifold Approximation and Projection," *Journal of Open Source Software*, vol. 3, no. 29, p. 861, 2018.
74. S. P. Ong, W. D. Richards, A. Jain, G. Hautier, M. Kocher, S. Cholia, D. Gunter, V. L. Chevrier, K. A. Persson, and G. Ceder, "Python Materials Genomics (pymatgen): A robust, open-source python library for materials analysis," *Computational Materials Science*, vol. 68, pp. 314–319, 2013.
75. C. J. Hargreaves, M. S. Dyer, M. W. Gaultois, V. A. Kurlin, and M. J. Rosseinsky, "The Earth Mover's Distance as a Metric for the Space of Inorganic Compositions," *Chemistry of Materials*, vol. 32, no. 24, pp. 10610–10620, 2020.
76. A. M. Glaudell, J. E. Cochran, S. N. Patel, and M. L. Chabinyk, "Impact of the Doping Method on Conductivity and Thermopower in Semiconducting Polythiophenes," *Advanced Energy Materials*, vol. 5, no. 4, p. 1401072, 2015.
77. S. B. Zhang, "The microscopic origin of the doping limits in semiconductors and wide-gap materials and recent developments in overcoming these limits: a review," *Journal of Physics: Condensed Matter*, vol. 14, no. 34, pp. R881–R903, 2002.
78. L. Sheng, L. Wang, T. Xi, Y. Zheng, and H. Ye, "Microstructure, precipitates and compressive properties of various holmium doped NiAl/Cr(Mo,Hf) eutectic alloys," *Materials & Design*, vol. 32, no. 10, pp. 4810–4817, 2011.
79. A. Mansouri Tehrani, A. O. Olynyk, Z. Rizvi, S. Lotfi, M. Parry, T. D. Sparks, and J. Brgoch, "Atomic Substitution to Balance Hardness, Ductility, and Sustainability in Molybdenum Tungsten Borocarbide," *Chemistry of Materials*, vol. 31, no. 18, pp. 7696–7703, 2019.
80. Mihailovich and Parpia, "Low temperature mechanical properties of boron-doped silicon," *Physical Review Letters*, vol. 68, no. 20, pp. 3052–3055, 1992.
81. Z. Qu, T. D. Sparks, W. Pan, and D. R. Clarke, "Thermal conductivity of the gadolinium calcium silicate apatites: Effect of different point defect types," *Acta Materialia*, vol. 59, no. 10, pp. 3841–3850, 2011.
82. T. D. Sparks, P. A. Fuierer, and D. R. Clarke, "Anisotropic Thermal Diffusivity and Conductivity of La-Doped Strontium Niobate Sr<sub>2</sub>Nb<sub>2</sub>O<sub>7</sub>," *Journal of the American Ceramic Society*, vol. 93, no. 4, pp. 1136–1141, 2010.
83. G. Grimvall, *Thermophysical Properties of Materials*. Amsterdam: North Holland, 1 ed., 1999.
84. R. Gaumé, B. Viana, D. Vivien, J.-P. Roger, and D. Fournier, "A simple model for the prediction of thermal conductivity in pure and doped insulating crystals," *Applied Physics Letters*, vol. 83, no. 7, pp. 1355–1357, 2003.
85. K. Pearson, "On lines and planes of closest fit to systems of points in space," *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, vol. 2, no. 11, pp. 559–572, 1901.
86. L. van der Maaten and G. Hinton, "Visualizing Data using t-SNE," *Journal of Machine Learning Research*, vol. 9, pp. 2579–2605, 2008.