PCCP



View Article Online **PAPER**



Cite this: Phys. Chem. Chem. Phys., 2025, 27, 22240

Received 10th June 2025 Accepted 22nd September 2025

DOI: 10.1039/d5cp02208h

rsc.li/pccp

Explainable GNN-derived structure-property relationships in interstitial-alloy materials

Eduardo Aquilar-Bejarano, (10 abc Luis Arrieta, (10 d Mauricio Gutiérrez, (10 e Ender Özcan, © ^c Simon Woodward, © ^{ab} Grazziela Figueredo © * f and J. Ignacio Borge (1)*9

This study presents a novel approach to understanding the structure-property relationships in nonstoichiometric materials and interstitial alloys using graph neural networks (GNNs). Specifically, we apply the crystal graph convolutional network (CGCNet) to predict the properties of transition-metal carbides, Mo₂C and Ti₂C, and introduce the crystal graph explainer (CGExplainer) enabling model interpretability. CGCNet outperforms traditional human-derived interatomic potential models (IAPs) in prediction accuracy and data efficiency, with significant improvements in the ability to extrapolate properties to larger supercells. Additionally, the CGExplainer tool enables detailed analysis of the relative spatial positioning of atomic ensembles, revealing key atomic arrangements that govern material properties. This work highlights the potential of GNN-based approaches for rapidly discovering complex structureproperty relationships and accelerating the design of materials with customized properties, particularly for alloys with variable atomic compositions. Our methodology offers a robust framework for future materials discovery, extending the applicability of GNNs to a broader range of materials systems.

1 Introduction

The relationship between atomic arrangements and material properties plays a fundamental role in materials science. This is particularly evident in two key types of metal alloys. The first type, "interstitial alloys", incorporate additional atoms (H, C, N, or B) into the existing interstitial spaces within the infinite crystal lattice of the metal while preserving the metallic lattice structure. Optimal material performance requires identification of optimal interstitial arrangements. 1-3 A significant sub-category within interstitial alloys are 'non-stoichiometric compounds', where the element proportions cannot be expressed as ratios of small natural numbers due to missing atoms, resulting in lattice vacancies. 4-6 The distribution of both interstitial atoms and

vacancies within the crystal lattice significantly influences various material properties, including mechanical strength, magnetic susceptibility, electrical resistivity,8 and catalytic properties.9-11 This influence stems from the pseudo-random distribution of interstitial atoms among available sites, creating numerous possible solid-state configurations, each exhibiting slightly different properties. 12-15

A second material type, substitutional alloys, demonstrate even more dramatically how atomic positioning influences material properties, particularly in emerging materials like two-dimensional (2D) transition metal carbides, nitrides, and carbonitrides (MXenes). These materials, with the general formula $M_{n+1}X_nT_x$ (where M represents early transition metals like Nb, V, or Ti; X is C and/or N where T_r denotes surface functional groups such as O, OH, or F; and n ranges from 1 to 4), 16-19 offer unprecedented opportunities for property manipulation through atomic arrangement. The ability to incorporate multiple transition metals at M sites creates solid-solution MXenes with enhanced functionalities, making the relationship between atomic arrangement and resulting properties particularly important. Recent research has demonstrated remarkable control over material properties can be achieved through selective metal substitution. For instance, Wang et al. showed that electrochemical properties in $Ti_{\nu}Nb_{2-\nu}CT_{x}$ and $V_{\nu}Nb_{2-\nu}CT_{x}$ can be systematically tuned by adjusting Ti: Nb and V: Nb ratios.20 Yang et al. established a strong correlation between Ti content and both electrical conductivity and carrier mobility in Ti_νNb_{2-ν}CT_x

^a GSK Carbon Neutral Laboratories for Sustainable Chemistry, University of Nottingham, Jubilee Campus, Triumph Road, Nottingham NG7 2TU, UK

^b School of Chemistry, The University of Nottingham, University Park, Nottingham NG7 2RD, Nottinghamshire, UK

^c School of Computer Science, The University of Nottingham, Jubilee Campus, Nottingham NG8 1BB, Nottinghamshire, UK

^d School of Chemical Engineering, Universidad de Costa Rica, San Pedro, San José, 11801 San José, Costa Rica

e School of Chemistry, Universidad de Costa Rica, San Pedro, San José, 11801 San José, Costa Rica

^f School of Medicine, The University of Nottingham, University Park, Nottingham NG7 2RD, Nottinghamshire, UK. E-mail: g.figueredo@nottingham.ac.uk

g Department of Physics, University of Alabama at Birmingham, Birmingham, AL 35294, USA. E-mail: jborgedu@uab.edu

Paper PCCP

MXenes.21 Additionally, Borge-Durán et al. predicted the emergence of ferromagnetic behavior in (Nb_{1-x}Ti_x)₄C₃T_x solid solutions, despite the non-magnetic nature of their end members, opening new possibilities for spintronic applications.²² These findings collectively demonstrate how precise control over atomic composition and positioning enables the rational design of materials with tailored characteristics.

To address the complex relationship between atomic arrangements and material properties, researchers developed empirical interatomic potential models (IAPs) as a systematic approach to understanding these correlations. IAPs work by identifying and characterizing specific atomic ensembles that contribute to desired material properties. Various approaches have been explored, from combinations of embedded atom method formalism with two-body Lennard-Jones interactions and three-body Axelrod-Teller potentials for Ti₂C systems, to more recent developments like the work of Borge-Durán et al., who successfully developed a potential for non-stoichiometric Mo₂C and Ti₂C.²³ However, the development of these potentials presents significant challenges. The process requires meticulous human expert identified atomic ensembles that contribute to the material's potential energy. While this task remains within human capabilities, it faces inherent limitations. The human mind, despite its analytical power, often struggles to perceive all subtle patterns and relationships within complex atomic arrangements. Furthermore, the vast number of possible atomic configurations creates an enormous search space. Consequently, the development of a single effective potential can extend from months to years, making this traditional approach increasingly impractical for rapid materials development.

Computational approaches based on graph neural networks (GNNs) offer potential solutions to these analytical limitations.²⁴ These advanced algorithms have shown remarkable success in decoding structure-activity relationships in various materials systems²⁵⁻³² by efficiently identifying correlations between graph-based structural representations and target properties. 33,34 However, despite their proven capabilities in various materials science applications, the potential of GNNs for understanding and predicting how relative atomic positioning modulates material properties remains largely unexplored compared to materials with well-defined structures and stoichiometry.

Existing GNN explainability methods focus primarily on understanding the relationship between material composition and properties through various analytical approaches. One notable implementation appears in crystal graph convolutional neural networks, where the separation of convolutional and pooling layers35 enables analysis of how specific atom types in given positions influence model predictions. Additional methodologies for explaining GNN-based materials property predictions have emerged in recent studies.³⁶⁻³⁸ However, these current approaches are fundamentally limited by their assumption of fixed atomic arrangements, analyzing only the impact of atom types at predetermined positions. This limitation becomes particularly significant when studying non-stoichiometric materials and solid solutions, where atoms distribute pseudo-randomly throughout the crystal lattice. A critical gap exists in current methodologies.

No method has been developed to explain models based on the relative three-dimensional positioning of atoms within crystal lattices. Addressing this gap is essential for understanding structure-property relationships in materials with complex atomic arrangements, particularly those where atomic positioning directly modulates material properties.

In this work, we evaluate the crystal graph convolutional neural network (CGCNet), 35 a specialized GNN architecture designed to predict the properties of materials, in the context of non-stoichiometric materials, with Mo₂C serving as our primary case study with further validation using Ti₂C. Both Mo₂C and Ti₂C belongs to the family of transition-metal carbides, interstitial alloys in which carbon atoms occupy interstitial spaces within a metallic lattice defined by metal atoms. This material system was selected due to its recent significant importance in materials science research39-44 and its structural similarity to MXenes, suggesting the potential transferability of our methodological approach. To complement CGCNet, we developed the crystal graph explainer (CGExplainer), a new tool that quantifies the contribution of specific atomic subassemblies and their relative spatial positions to material properties, enabling systematic analysis of structure-property relationships in threedimensional space. Our investigation addresses three fundamental objectives: (1) assessment of CGCNet capabilities in predicting properties of novel atomic configurations within solid solutions; (2) evaluation of GNN performance in extrapolating properties for larger supercells with previously unobserved atomic configurations; and (3) demonstration of CGExplainer's ability to interpret GNN models trained on crystal graph data, revealing correlations between atomic ensembles and target properties.

2 Methods

We aim to accurately predict the ground-state energies of transition-metal carbides using graph representations, targeting results close to those obtained from high-level DFT calculations. To achieve this goal, we have followed the workflow outlined in Fig. 1.

This workflow consisted of five main steps: (1) generation of the Mo_2C dataset by creating 1065 Mo_2C 2 × 2 × 2 structures and calculating their energy using DFT; (2) creation of a graph representation of these structures; (3) training and evaluating the GNN model in comparison to the human-derived IAP for the Mo₂C and Ti₂C dataset; (4) assessment of the models' performance and ability to extrapolate using a dataset of $2 \times 2 \times 4$ Mo₂C structures; and (5) explanation of the CGCNet models using our newly proposed method, the crystal graph explainer (CGExplainer), and comparison of its explanations between the Mo₂C and Ti₂C dataset.

2.1 The Mo₂C and Ti₂C dataset

The Mo₂C dataset was built from scratch, starting with a $1 \times 1 \times 1$ cubic cell of MoC as the foundation for constructing 2 \times 2 \times 2 supercells with a 1:1 stoichiometry. To achieve a 2:1 stoichiometry, half of the carbon atoms were removed from the supercells.

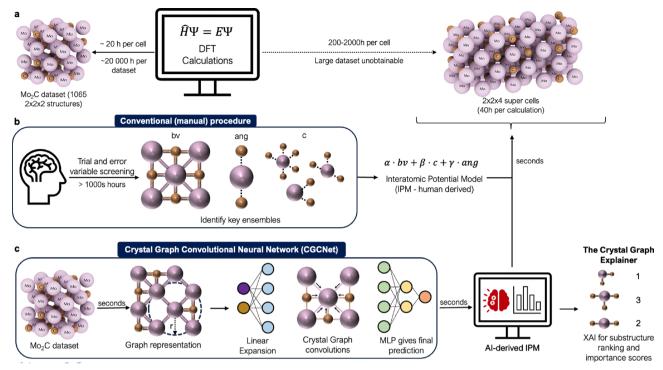


Fig. 1 Procedure followed herein. (a) Creation of the Mo_2C dataset and computational cost of calculation of energy for $2 \times 2 \times 4$ supercells. (b) Conventional paradigm to create atomic potentials based on atomic ensembles discovered by human trial and error. (c) Automatically GNN created interatomic potentials for non-stoichiometric materials and further understanding of those contributing atomic ensembles by using the crystal graph explainer.

The special quasi-random structures (SQS) algorithm⁴⁵ was employed to reproduce the local atomic disorder arrangements typically found in random solid solutions. This method generated a diverse collection of carbon/vacancy configurations within the Mo₂C supercells, effectively emulating the varied atomic arrangements present in real materials. Further details are provided in SI under 'generating C/vacancy configurations in Mo₂C supercells for database creation'.

To investigate the energetics of the diverse cubic Mo₂C configurations, we performed first-principles calculations using density functional theory (DFT) as implemented in the quantum ESPRESSO package. 46 The exchange-correlation energy was computed using the Perdew-Burke-Ernzerhof (PBE) functional⁴⁷ within the generalized gradient approximation (GGA) framework. Ultra-soft pseudopotentials from the Rutgers University GBRV database⁴⁸ were employed to describe the interactions between valence and core electrons. The Kohn-Sham orbitals were expanded using a plane-wave basis set with an energy cutoff of 40 Ry. Brillouin zone integration was performed using a $4 \times 4 \times 4$ Monkhorst-Pack⁴⁹ k-point grid.

In the case of the Ti₂C dataset, this was taken from the work by Borge-Durán et al.23

2.2 Graph representation

To generate the graph representation that serves as input to the GNN, we used the DFT-relaxed structures of Mo₂C and Ti₂C. In our graph representation, each node represents an atom, while edges represent atom-atom interactions. To define the

connectivity between atoms, we employed a cut-off radius centered on each atom, where all atoms within this boundary are considered to share an edge with the central atom (see Fig. S1a in SI). To accurately model the infinite lattice of the material, we applied periodic boundary conditions (PBC) to the graph. This ensures that nodes located on opposite faces of the crystal are appropriately connected (see Fig. S1b in SI).

As node features, we used the chemical element each node represents, one-hot-encoded. As edge features, we used the solid angle, defined in eqn (1), where r is the distance between an atom and a surface and A is the area of such surface. The surfaces result from the faces of the polyhedrons were created by applying a Voronoi tessellation algorithm^{50,51} considering each atom as a point in space (see Fig. S1c and d in SI). The solid angle is an ideal descriptor, as this quantity changes depending on the relative positioning of carbide atoms around molybdenum atoms (see Fig. S2 and Crystal graph section in SI). We used a Gaussian expansion to increase the dimensionality of the solid angle (eqn (1)) to improve the performance of the model, as done before by Gu et al. to model perovskites.30

$$\Omega = \frac{A}{r^2} \tag{1}$$

To train the GNN, each graph has been labeled with it's corresponding DFT ground state energy. In the case of the transition metal carbides, the most stable structure is known (i.e. lowest energy structure). We used the energy of this structure as baseline (E = 0) and predicted the difference of

energy between the given structure and this baseline. This decision was made as usually this calculation is usually performed to quantify differences in energy between structures rather than to attain absolute energies. Our final target variable is therefore defined as follows:

$$E_{\text{target}} = E_{\text{DFT}} - E_{\text{DFT}}^{0} \tag{2}$$

Model architecture

GNNs are a type of artificial neural network that can be trained on graphs, which allows understanding of relations of graph structured data and outcome variables. The CGCNet consists of two phases: (1) message passing and (2) readout. The first phase is a node-level operation block, which explores the topology of the graph to capture the complex relations between neighboring nodes. This operation is known as convolution. The convolutional operator used for this network is the 'crystal graph convolution', 35 which is defined in eqn (3), where x_i is the node feature vector of the central node i, $z_{i,i} = [x_i, x_i, e_{i,i}]$ denotes the concatenation of features of the central node i, neighbor node j, and the edge between them, σ represents the sigmoid function, g represents the softplus function, and W_f , b_f , W_s , and $b_{\rm s}$ are learnable parameters.

$$\mathbf{x'}_{i} = \mathbf{x}_{i} + \sum_{j \in \mathcal{N}(i)} \sigma(\mathbf{z}_{i,j} \mathbf{W}_{f} + \mathbf{b}_{f}) \odot g(\mathbf{z}_{i,j} \mathbf{W}_{s} + \mathbf{b}_{s})$$
(3)

The second phase is a prediction occurring at the graph level. Its first step is to summarize the information from all the nodes contained within the graph into a single graph-level feature vector. This is achieved by an operation called pooling. This operation is usually an element-wise operation that runs for all the node's feature vectors contained within the graph. Pooling ensures that all the graphs within the database attain a single fixed-size vector representing them. As we are predicting an extensive property and the models are all trained with a fixed size supercell $(2 \times 2 \times 2)$, we adopted a sum pooling operation. Using another type of pooling, such as max or average, would deliver predictions within the range of values of training data, regardless of the number of nodes (atoms) in the graph. In sum pooling, the GNN is effectively learning the energetic contribution of each atom, thus making the model aware of the relation between energy and quantity of atoms. Finally, this latter vector is the input into a multilayer perceptron (MLP) that outputs a final prediction. Our GNN steps from graph input to prediction are:

- The node features are taken (length of two) and are expanded to a final length of 16 by a fully connected layer with no activation function, $R^{\text{nodes}\times 2} \rightarrow R^{\text{nodes}\times 16}$
- The crystal graph convolution operator updates the node states in of all nodes a total of two times, $R^{\text{nodes} \times 16} \rightarrow R^{\text{nodes} \times 16}$.
- Sum pooling is applied to all the node feature vectors to get a graph-level feature vector, $R^{\text{nodes} \times 16} \rightarrow R^{1 \times 16}$.
- A fully connected layer with leakyReLU activation function takes the graph-level vector and maps it to half of its original length, $R^{1\times16} \rightarrow R^{1\times8}$.
- · A last fully connected layer with no activation function transforms the feature vector into a single scalar number, this being the prediction of the model, $R^{1\times 8} \rightarrow R$.

2.4 Crystal graph explainer

Human-derived interatomic potentials have the advantage of being highly interpretable, while GNNs usually are not.

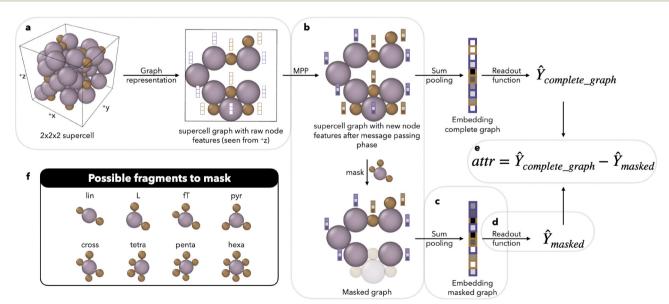


Fig. 2 Crystal graph explainer algorithm. (a) The algorithm takes an input graph representation of a material and fits it into a pre-trained GNN model. (b) Before the sum pooling, the algorithm creates a copy of the graph with the embeddings obtained during the message passing phase (MPP in the figure). (b) The algorithm then masks the substructure which an attribution score is to be assigned to, and (c) then the pooling operation occurs. (d) Lastly, the new graph embedding is fitted into the readout phase which leads to a prediction. (e) We define the attribution score of such fragment as the difference between the prediction of the complete graph and the prediction of the masked graph, as done before by Wu et al. 52 (f) Shows the fragments that is possible to mask in the Mo₂C structures

To overcome such a problem, we designed a unique crystal graph explainer (CGExplainer). The CGExplainer is a graph perturbation method designed to explain GNN models that predict properties of solid-state crystal structures. As proposed by Wu et al. for organic molecules,52 the algorithm works by masking certain fragments of the structure just before the pooling operation. A scheme of this approach is shown in Fig. 2.

As the GNN (CGCNet, in our case) learns during training, it identifies the most important nodes required for making accurate predictions. In other words, if the network recognizes that an ensemble of atoms plays an important role in determining the target property, the pooling operation will prioritize gathering information from those atoms.

Consequently, if a specific atomic ensemble is masked, the resulting graph embedding will likely differ from that of the complete graph—provided those atoms were indeed important for the prediction (see the SI, Crystal graph explainer section, for further details). This difference in embeddings leads to a change in the model's prediction.

We define the attribution (importance) of a masked fragment as the difference between the prediction made with the complete graph and the prediction made with the masked graph. As the perturbation made to get an attribution score is done at graph level, the methodology presented is model agnostic, meaning that CGExplainer can be applied to any trained GNN model.

When attempting to mask a substructure in a crystal graph, several problems may arise because of the high degree of repetitive patterns throughout the cell. The issues we found and the proposed solutions are discussed in SI Crystal graph explainer section.

2.5 The interatomic potential benchmark

As benchmark, the interatomic potential by Borge-Durán et al.²³ was used. This potential consists of a multiple linear regression fitted to the internal energy of Mo₂C or Ti₂C materials, which consisted of three parameters. This potential was chosen because it was specifically designed to describe the energetics of the materials studied herein, thus being the most accurate model found in the literature. Notably, this IAP was previously validated against experiment by reproducing order-disorder transition temperatures for Mo₂C and Ti₂C (pred. 1430/1430 °C; exp. 1030/1024 °C), see ref. 23.

The first parameter is the bond valence parameter. This parameter responds to the fact that, in a stable structure, the total valence of each atom will be equal to or close to its preferred valence. In the case of molybdenum, the preferred valence is 6, and for titanium it's 4. The deviation of this ideal behavior was quantified by eqn (4), where n is the total number of transition metal atoms in the structure, i represents the i-th transition metal atom in the structure, V_i is the valence of the atom i, and V_0 is the preferred valence of of the atom i.²³

$$E_{\rm bv} = \sqrt{\frac{1}{n-1} \cdot \sum_{i} (V_i - V_0)^2}$$
 (4)

The second parameter quantifies how homogeneous the carbon atoms are distributed within the lattice. The rationale behind this parameter is that a stable structure will favor a homogeneous distribution of atoms (3 carbon atoms per metal atom) over conglomerates of carbides. This term is calculated as shown in eqn (5), where the number of carbons around the i-th metal atom is calculated, and then elevated to the power of two. The resulting parameter consists of the sum of this number for all metal atoms in the system.

$$E_{\rm C^2} = \sum_{\rm Mo} \left(\sum_i C_i \right)^2 \tag{5}$$

The last parameter is an empiric parameter related to the distribution of carbide atoms in the lattice. It was found that structures with numerous C-Mo-C at a 180° arrangement exhibit higher energies than those structures that did not have such structure. This way, the last parameter is calculated as shown in eqn (6).

$$E_{180^{\circ}} = \sum_{i} N_{180^{\circ}} \tag{6}$$

The resulting potential consisted in a multiple linear regression of these three parameters, as shown in eqn (7)

$$E = \alpha \cdot E_{\text{bv}} + \beta \cdot E_{\text{C}^2} + \gamma \cdot E_{180^{\circ}} \tag{7}$$

3 Results and discussion

Model performance for the $2 \times 2 \times 2$ supercells

To evaluate both approaches robustness and ability to generalize, we applied a nested cross validation approach (further details in section Model training and Fig. S3 in SI). We created five-folds, which led to a total of five test sets, each being evaluated by four different training and validation sets. This generated a total of 20 training test processes. This exhaustive splitting strategy reduces the possibility of testing the models at a set of points where they perform particularly well. Both IAP and CGCNet were trained and evaluated using the same set of folds. For CGCNet, the training set was used to learn weights and biases and the validation set to update the learning rate and early stop. For IAP, the training set was used to determine the linear regression coefficients; the validation set had no impact on the training for this method. The test sets provided an unbiased evaluation of the performance of both models.

To present the results, we aggregated the predictions from all outer test folds of the nested cross-validation. For each data point, we calculated the mean predicted value across folds. These mean predictions are shown in the parity plot in Fig. 3, for both Mo₂C and Ti₂C datasets, and Table 1 shows the performance metrics for both datasets. A detailed analysis of the performance in each individual test fold is provided in the SI in Fig. S5 and S6.

From the metrics in Table 1, it is clear that CGCNet performs better in the predictive tasks than the human expert derived IAP of Borge-Durán et al.23 To test this, we applied a Wilcoxon test and found that there is statistical difference between CGCNet

Table 1 Predictive performance of CGCNet and IAP for 2 \times 2 \times 2 supercells. Reported are the coefficient of determination (R^2) , mean absolute error (MAE), and root-mean-square error (RMSE) obtained for all the testing points from the nested-CV; higher R² and lower MAE/RMSE indicate better accuracy

	Mo ₂ C			Ti ₂ C		
Model	R^2	MAE (eV)	RMSE (eV)	R^2	MAE (eV)	RMSE (eV)
CGCNet IAP		0.087 0.126	0.115 0.160		0.056 0.089	0.109 0.146

and IAP for both Mo_2C (p-value = 0.036) and Ti_2C (p-value = 4.77×10^{-5}) datasets. This behavior is expected as Borge-Durán's potential consists of only three variables which are combined linearly, while the CGCNet model consists of hundreds of learnable parameters that adjust in a data-driven way and are combined non-linearly, better modeling the real behavior of the material.

From the parity plots in Fig. 3, it is noticeable that the energy values of the two materials are very different, ranging from 0 to 5 eV for Mo₂C and from 0 to 3 eV for the case of Ti₂C. Also, in the case of the Mo₂C, the DFT energy values are well populated in the whole energy range, while in the case of Ti₂C there is a gap from 0 to 2 eV and from 2.7 to 3 eV, suggesting that the latter is a less diverse dataset than the former. The reason for these differences is that the datasets were created using random generators. As the Mo₂C dataset is more populated (1065 structures), the probability of having more diverse structures (and therefore better sampled energy values) is higher compared to the Ti₂C dataset. As the Mo₂C is well sampled, it is expected that the trained models, for IAP and CGCNet, are robust in predicting the energy of potentially any unseen structure. In the case of models trained on the Ti₂C

Table 2 Predictive accuracy of CGCNet versus IAP on Mo₂C supercells. Listed are test-set means for R^2 , MAE, and RMSE; higher R^2 and lower MAE/ RMSE denote better performance

Model	$1 \times 1 \times 2$			2 imes 2 imes 4		
	R^2	MAE (eV)	RMSE (eV)	R^2	MAE (eV)	RMSE (eV)
CGCNet IAP		0.085 0.242	0.172 0.297		0.697 1.617	0.904 1.831

dataset, it is expected that the generalization is more limited, as the models have seen fewer material configurations. However, it is reassuring the fact that both IAP and CGCNet are accurate predicting the least and most stable structure (higher and lower energy value, respectively), demonstrating that even in extreme cases where the models have not seen structures with similar values of energy, they are still able to deliver acceptable predictions, indicating that it is likely that materials with energies in the whole range from 0 to 3 eV will have an accurate prediction.

To assess the data efficiency of the models, we created a learning curve for the Mo₂C dataset. This plot shows the performance of the model when changing the size of the dataset. The results are shown in Fig. 4, where each point represents the mean absolute error through the 20 train test processes and the error bars the standard deviation (plots showing the performance based on RMSE and R^2 is shown in Fig. S5 in SI). We found that when the total size of the Mo₂C dataset is reduced to only 300 structures and split in a 3:1:1 ratio, CGCNet outperforms IAP. This suggests that our methodology still achieves a satisfactory result with 180 data points for training, demonstrating that the workflow is data efficient.

Although CGCNet only needs 300 structures to achieve comparable performance to the IAP, the latter exhibits more

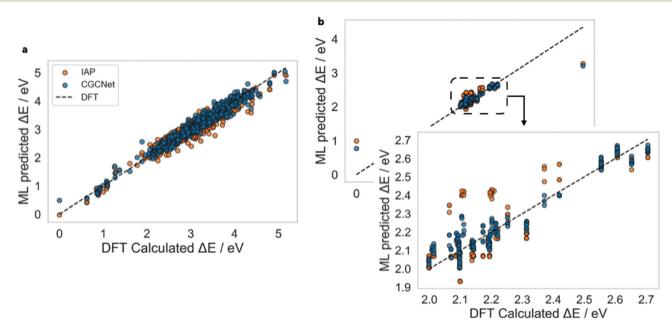


Fig. 3 Parity plot of the predictions of internal energy of transition metal carbides using the IAP and CGCNet methods. (a) Shows the parity plot for Mo₂C and (b) shows the parity plot for Ti₂C.

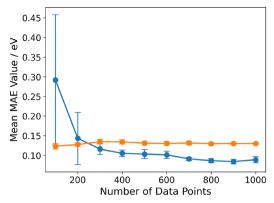


Fig. 4 Learning curve for CGCNet and IAP. The y-axis shows the mean absolute error, while the x-axis shows the total quantity of points in the

stable performance regardless of the number of points in the database (Fig. 4). This behavior is expected, as the IAP consists only of three variables fitted into a multiple linear regression. As only three parameters are learned in the training process, just a few structures are necessary to optimize them. Therefore, even from a small number of structures, very accurate models can be achieved using the IAP. This demonstrates that, in case of modeling databases with a limited number of structures, the IAP still presents the advantage of delivering accurate predictions. However, when mid-size databases are modeled (from 300 or more), CGCNet can be used to get models at least as accurate as the human-derived IAP.

3.2 Generalizing energy prediction models to supercells of unseen sizes

The ultimate goal of atomic potentials is to predict properties of unstudied atomic configurations or accurately model larger systems where DFT calculations are computationally expensive. To test this here, we predict the internal energy of $1 \times 1 \times 2$ and $2 \times 2 \times 4$ Mo₂C supercells. Although DFT calculations for

 $1 \times 1 \times 2$ supercells can be performed, we included this as a final test set to evaluate the ability of both approaches to extrapolate to both larger and smaller supercells. The internal energy of the most stable supercell for each size was calculated and set as the new reference point. As the energy of these new, more stable cells differs in magnitude from the energy of the most stable cell on which the models were trained, it was necessary to normalize both the IAP and CGCNet predictions. Therefore, we define our target variable as the difference between the energy of the ith cell and the energy of the most stable Mo₂C supercell of the same size. Since 20 models arose from the nested cross-validation approach, we adopted an averaging system where each model predicted the energy difference of each cell, and the reported result was the mean of the predictions from the different models. The results are shown in Fig. 5, and the metrics obtained are presented in Table 2.

Fig. 5a shows the parity plot for the $1 \times 1 \times 2$ supercells. Both models deliver predictions with satisfactory accuracy; however, the metrics in Table 2 show that CGCNet achieves higher accuracy with statistical difference (p-value = 0.002). For the case of $2 \times 2 \times 4$ supercells (Fig. 5b), both models tend to underestimate the energy difference; however, the humanderived IAP predicts it less accurately, as shown in Table 2. We also found statistical difference for this dataset as well (p-value = 0.011). The reason for the poorer performance of the IAP may be that the prediction scaling grows linearly with the size of the cell, which does not represent the nature of the material. An additional limitation of IAP is that it is based only on ensembles that consider direct neighbors of atoms. For GNNs, although our graph representation considers only direct neighbors, the convolutions in the message-passing phase allow the model to capture long-range interactions, which are ignored in IAP. While our experiment does not demonstrate the practical application of these extrapolations to real-life scenarios (i.e. predict more complex properties such as band gaps), the proof of concept demonstrates that such modeling is possible and should be explored in future studies.

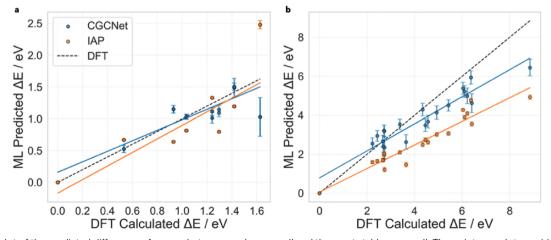


Fig. 5 Parity plot of the predicted difference of energy between each supercell and the most stable supercell. The points are plot considering the mean of the prediction of all the models trained in Section 3.1 and the error bars represent the standard deviation. (a) Shows the results of the $1 \times 1 \times 2$ supercell test set and (b) shows the results of the $2 \times 2 \times 4$ supercell test set.

3.3 Explainability of the CGCNet models

While Fig. 3 demonstrates that CGCNet outperforms the IAP, the latter has the advantage of being highly interpretable. We used our CGExplainer to rationalize which relative position in the atomic arrangements modulates the properties of the interstitial alloys studied herein (both Mo₂C and Ti₂C). To perform these analyses, a random model from the 20 developed was taken and the analysis was performed using the test set for all atom ensembles shown in Fig. 2f. The attribution scores of the CGExplainer for each material were normalized between 0 and 1 to allow comparison, where values closer to 0 are assigned to ensembles with lower attributions (lower impact on the target property). The results are shown in Fig. 6.

Our graph-neural-network evaluation of local coordination environments shows that both Mo₂C and Ti₂C strongly favour the trigonal-pyramidal ('pyr') motif: it attains the lowest normalised attribution score (0.00), indicating the highest stability in both carbides.

The lowest-energy arrangement in both Mo₂C and Ti₂C is a trigonal-pyramidal coordination (C_3 symmetry) in which the three carbon ligands form an equilateral base with C-C-C angles of $\approx 120^{\circ}$, while one metal s orbital mixes with two d orbitals to create three equivalent sd² hybrids that point directly toward the carbon atoms and maximise M-C σ overlap.⁵³ Although mutually orthogonal sd2 hybrids would impose ideal C-M-C angles of 90°, ligand-ligand repulsion within the extended lattice distorts them toward $\sim 120^{\circ}$, a characteristic feature of many 'non-VSEPR' d⁰ transition-metal compounds.⁵⁴ Motifs that cannot sustain this hybridisation (e.g. flat-T or linear) display weaker σ interactions and are thus higher in energy. Despite formal oxidation states approaching +4, M-C bonding remains partly covalent and partly metallic, with substantial electron sharing and delocalisation rather than full charge transfer.55 Residual (non-hybridised) d orbitals participate in secondary interactions or remain non-bonding, further stabilising the electronic structure, and the preference for the pyramidal motif is accentuated in Mo₂C because the more diffuse 4d orbitals of molybdenum enhance medium-range overlap relative to the compact 3d orbitals of titanium, which favour tighter coordinations.⁵⁶

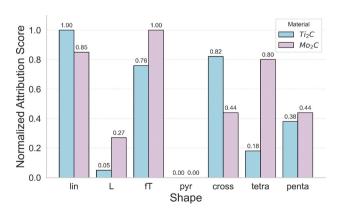


Fig. 6 Comparison of importance of atomic ensembles found in the nonstoichiometric materials Mo₂C and Ti₂C estimated using the CGExplainer (see Fig. 2f for ensemble abbreviations)

The secondary coordination preferences of Mo₂C and Ti₂C are governed by intrinsic differences between 4d and 3d valence shells. In Mo₂C, geometries ranked after the pyramidal minimum (0.00) are the 'L' (0.27) and 'cross' (0.44) motifs, whereas Ti₂C stabilises 'L' (0.05) and then 'tetra' (0.18; see Fig. 2 for ensemble labels). These trends arise because 4d orbitals possess an additional radial node and are more spatially diffuse than their 3d counterparts, enlarging the radial extent of Mo valence density and enabling longer-range metal-carbon overlap.⁵⁷ Conversely, the compact 3d set of Ti enhances short-range σ interactions, favouring tighter coordinations such as the tetrahedral environment. The much higher attribution score of 'tetra' in Mo₂C (0.80) relative to Ti₂C underscores how orbital diffuseness modulates relative stability.

The least favourable motifs in both carbides, namely, 'fT' (flat-T) and 'lin' (linear), induce d-orbital splitting patterns poorly matched to either 3d or 4d electron distributions, leading to pronounced destabilisation. Overall, the CGExplainer attribution hierarchy mirrors established electronic-structure principles: the greater radial extension and lower pairing energies of second-row (4d) metals shift their coordination preferences away from the compact geometries preferred by first-row (3d) analogues.⁵⁸ This consistency between machine-learning explanations and ligand-field expectations validates the physical soundness of the model.

As the importance of atomic ensembles is estimated by CGExplainer, these could potentially be used as a guide to create traditional IAPs. For example, from Fig. 6, it is clear that fragments containing carbon-metal-carbon subunits in a 180° position are considered to contribute more to the energetics of the material than those fragments with the same number of carbide atoms without such conformation (see the attributions of "lin" compared to "L" and the attributions of "fT" compared to "pyr"). This observation is in agreement with the empiric discovery of Borge-Durán et al. that carbide atoms in a 180° conformation lead to higher energy materials.23

Beyond guiding IAP design, this agreement also serves as an experimentally anchored validation of what CGCNet learns. In our prior work, 23 a human-designed IAP based on three motifs, namely, a low-energy trigonal-pyramidal C-M-C units, penalized 180° C-M-C arrangements, and an approximately homogeneous carbon distribution captured by E_{C^2} , reproduced order-disorder transition temperatures for Mo₂C (pred. 1430 °C; exp. 1430 °C) and Ti_2C (pred. 1024 °C; exp. 1030 °C). Here, CGCNet + CGExplainer independently recovers the same rules from the DFT data (low attribution for pyramidal C-M-C, high for 180° C-M-C, and consistency with the E_{C^2} trend), indicating that two independent approaches converge on motifs tied to measurable observables. Thus, the IAP comparison is not merely a computational baseline but a bridge to experiment that supports the physical validity of the learned representations.

However, the derivation of more complex variables that explain the energetics of the material are not as easy to obtain. For example, from Fig. 6, there is no trend between the importance of the fragment and the total number of carbide

atoms in such fragment, making the derivation of the E_{C^2} parameter (eqn (5)) nearly impossible from the explanations. This way, while the estimated importance of ensembles may help guiding the selection of variables for the construction of traditional IAPs, the derivation of a complete potential from these estimations is very difficult, and should be complemented by knowledge of the system to attain complete and accurate traditional potentials.

4 Conclusions

GNN models present a significant advantage over conventional IAPs as they allow the finding of key atomic ensembles automatically. Although traditional IAPs have been widely used in materials sciences, 59-62 their construction is slowed by the complexity and the large number of possible key ensembles that may correlate to a property of interest, which arises a very challenging task for humans experts. GNNs can find such ensembles from a graph representation of a material efficiently and rapidly. Particularly, CGCNet models, compromising the GNN architecture and crystal graph representation, are powerful new tools affording property predictions more rapidly than the best human experts.

A typical limitation identified in GNNs is that it lacks interpretability, but is overcome with the CGExplainer. This tool allows assigning an attribute score to a given fragment contained within the material. The tool shows excellent agreement with known behavior of the materials. This demonstrates that the GNN is able to rank important lattice fragments similarly or better than human experts. The differences found in calculated importance between CGExplainer and human methodologies can be the reason why the IAP performs worst than the GNN. Furthermore, the CGExplainer tool allows rapid analysis of general trends in quantitative structure-property relationships for a wide range of materials considering the 3D disposition of atoms in the cell, which is not possible yet for current methods. Our method can especially be useful for interstitial alloys or materials where the atomic composition is of high variety, such as high entropy materials.

The prediction of energy in $2 \times 2 \times 4$ cells demonstrated that CGCNet allows for more accurate extrapolations to larger systems where DFT calculations are not yet computationally feasible. As the creation of IAPs is usually motivated by the lack of computational power to calculate properties of larger systems, our results demonstrated that GNNs are capable of predicting properties of unknown atomic configurations and those from larger cell sizes. The results obtained encourage the use of these algorithms in combination with CGExplainer to attain high accuracy potentials while enabling model interpretability.

Our workflow, combining the crystal graph network (CGCNet) and the crystal graph explainer (CGExplainer), shows great promise in defining structure-property relationships in non-stoichiometric materials and interstitial alloys. CGCNet outperformed the traditional human-derived interatomic potential model (IAP) in terms of accuracy and generalizability,

while CGExplainer provided interpretable explanations for the GNN model's predictions. The automated discovery of key atomic ensembles and the ability to explain their influence on material properties pave the way for a more efficient and insightful materials design process.

Future research will focus on applying this approach to a wider range of materials, such as MXenes and high-entropy alloys, to validate its generalizability and potential impact on materials science. By extending this methodology to diverse material systems, we can accelerate the development of novel materials with tailored properties and unlock new possibilities in the field of materials engineering.

5 Computational tools

CrysGNet was built on PyTorchGeometric 2.3.163 running over PyTorch 2.0.1.64 The graphs were built using PyTorchGeometric 2.3.163 and PyMatGen 2023.8.10.65 We have developed a modified version of the stratification and cross-validation procedure developed by Pablo-García et al.66 Masking procedure was developed exclusively for this work using PyTorchGeometric 2.3.1⁶³ and NetworkX 3.0.⁶⁷ All the plots were built using Matplotlib 3.7.3⁶⁸ and Seaborn 0.13.0.⁶⁹ Structure images were generated using Ovito⁷⁰ and Vesta.⁷¹

Author contributions

Conceptualization, E. A. B., and I. B.; methodology, E. A. B., E. Ö., S. W., G. F., and I. B.; software, E. A. B. and I. B.; investigation, E. A. B., L. A., M. G., E. Ö., S. W., G. F., and I. B.; resources, L. A., M. G., and I. B.; data curation. E. A. B., L. A., M. G., and I. B.; writing – original draft, E. A. B., and I. B.; writing - review and editing, E. Ö., S. W., G. F., and I. B.; supervision, E. Ö., S. W., G. F., and I. B.

Conflicts of interest

There are no conflicts to declare.

Data availability

Code and data for this article, including the Mo₂C and the Ti₂C databases are available at Ineterstitial Alloys Repository at https://github.com/EdAguilarB/InterstitialAlloys.

Supplementary information: supplementary figures and detailed procedures. Predictions and metrics for all the models trained in this study. Supplementary information is available. See DOI: https://doi.org/10.1039/d5cp02208h.

Acknowledgements

E. A. B. is grateful for financial support of the Artificial Intelligence Doctoral Training Centre of the School of Computer Sciences of the University of Nottingham.

Notes and references

Paper

- 1 C. Liu, W. Lu, W. Xia, C. Du, Z. Rao, J. P. Best, S. Brinckmann, J. Lu, B. Gault, G. Dehm, G. Wu, Z. Li and D. Raabe, *Nat. Commun.*, 2022, 13, 1102.
- 2 W. Pepperhoff and M. Acet, *Constitution and Magnetism of Iron and its Alloys*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2001, pp. 147–181.
- 3 T. Chen, C. Foo and S. C. Edman Tsang, *Chem. Sci.*, 2021, **12**, 517–532.
- 4 Y. Zhou, T. W. Heitmann, E. Bohannan, J. C. Schaeperkoetter, W. G. Fahrenholtz and G. E. Hilmas, *J. Am. Ceram. Soc.*, 2020, **103**, 2891–2898.
- 5 A. I. Gusev, A. A. Rempel and A. J. Magerl, Disorder and Order in Strongly Nonstoichiometric Compounds: Transition Metal Carbides, Nitrides and Oxides, Springer Berlin Heidelberg, Berlin, Heidelberg, 2001, pp. 43–112.
- 6 M. Yoshitake, J. Vac. Sci. Technol., A, 2014, 32, 061403.
- 7 S.-H. Jhi, S. G. Louie, M. L. Cohen and J. Ihm, *Phys. Rev. Lett.*, 2001, **86**, 3348.
- 8 W. S. Williams, *Int. J. Refract. Met. Hard Mater.*, 1999, 17, 21–26.
- C. B. Krishnamurthy, O. Lori, L. Elbaz and I. Grinberg,
 J. Phys. Chem. Lett., 2018, 9, 2229–2234.
- 10 X. Liu, C. Kunkel, P. Ramirez de la Piscina, N. Homs, F. Vines and F. Illas, ACS Catal., 2017, 7, 4323–4335.
- 11 N. Kuriakose, K. Mondal and P. Ghosh, et al., Phys. Chem. Chem. Phys., 2020, 22, 14599-14612.
- 12 I. Baker, B. Grabowski, S. V. Divinski, X. Zhang and Y. Ikeda, *MRS Bull.*, 2023, **48**, 769–776.
- 13 B. D. Tinh, N. Q. Hoc, D. Q. Vinh, T. D. Cuong and N. D. Hien, *Adv. Mater. Sci. Eng.*, 2018, **2018**, 5251741.
- 14 A. I. Gusev, A. A. Rempel and A. J. Magerl, Disorder and Order in Strongly Nonstoichiometric Compounds: Transition Metal Carbides, Nitrides and Oxides, Springer Berlin Heidelberg, Berlin, Heidelberg, 2001, pp. 299–331.
- 15 A. I. Gusev, A. A. Rempel and A. J. Magerl, Disorder and Order in Strongly Nonstoichiometric Compounds: Transition Metal Carbides, Nitrides and Oxides, Springer Berlin Heidelberg, Berlin, Heidelberg, 2001, pp. 453–601.
- 16 Y. Gogotsi and Q. Huang, ACS Nano, 2021, 15, 5775-5780.
- 17 Y. Gogotsi and B. Anasori, ACS Nano, 2019, 13, 8491-8494.
- 18 K. R. G. Lim, M. Shekhirev, B. C. Wyatt, B. Anasori, Y. Gogotsi and Z. W. Seh, *Nat. Synth.*, 2022, **1**, 601–614.
- 19 R. M. Ronchi, J. T. Arantes and S. F. Santos, *Ceram. Int.*, 2019, **45**, 18167–18188.
- 20 L. Wang, M. Han, C. E. Shuck, X. Wang and Y. Gogotsi, *Nano Energy*, 2021, **88**, 106308.
- 21 Y. Yang, M. Han, C. E. Shuck, R. K. Sah, J. R. Paudel, A. X. Gray, Y. Gogotsi and S. J. May, 2D Mater., 2022, 10, 014011.
- 22 I. Borge-Durán, A. Paul and I. Grinberg, *Chem. Mater.*, 2023, 35, 7442–7449.
- 23 I. Borge-Durán, D. Aias and I. Grinberg, *Phys. Chem. Chem. Phys.*, 2021, 23, 22305–22312.
- 24 P. Reiser, M. Neubert and A. Eberhard, et al., Commun. Mater., 2022, 3, 93.

- 25 K. Choudhary and B. DeCost, npj Comput. Mater., 2021, 7, 185.
- 26 S.-Y. Louis, et al., ACS Appl. Mater. Interfaces, 2022, 14, 26587–26594.
- 27 J.-Y. Ryu, E. Elala and J.-K. K. Rhee, *Materials*, 2023, 16, 4300.
- 28 J. N. Law, et al., JACS Au, 2023, 3, 113-123.
- 29 V. Fung, J. Zhang, E. Juarez and B. G. Sumpter, npj Comput. Mater., 2021, 7, 84.
- 30 G. H. Gu, J. Jang, J. Noh, A. Walsh and Y. Jung, npj Comput. Mater., 2022, 8, 71.
- 31 J. Zhang, A. Koneru, S. K. R. S. Sankaranarayanan and C. M. Lilley, ACS Appl. Mater. Interfaces, 2023, 15, 20520–20530.
- 32 S.-Y. Louis, Y. Zhao, A. Nasiri, X. Wang, Y. Song, F. Liu and J. Hu, *Phys. Chem. Chem. Phys.*, 2020, 22, 18141–18148.
- 33 J. Zhou, G. Cui, S. Hu, Z. Zhang, C. Yang, Z. Liu, L. Wang, C. Li and M. Sun, *AI Open*, 2020, 1, 57–81.
- 34 S. Zhang, H. Tong, J. Xu and R. Maciejewski, *Comput. Soc. Networks*, 2019, **6**, 11.
- 35 T. Xie and J. C. Grossman, *Phys. Rev. Lett.*, 2018, **120**, 145301.
- 36 J. Schmidt, L. Pettersson, C. Verdozzi, S. Botti and M. A. L. Marques, *Sci. Adv.*, 2021, 7, eabi7948.
- 37 D. Cheng, W. Sha, Q. Han, S. Tang, J. Zhong, J. Du, J. Tian and Y.-C. Cao, *Electrochim. Acta*, 2024, 473, 143459.
- 38 M. Dai, M. F. Demirel, Y. Liang and J.-M. Hu, *npj Comput. Mater.*, 2021, 7, 103.
- 39 Y. Deng, Y. Ge, M. Xu, Q. Yu, D. Xiao, S. Yao and D. Ma, *Acc. Chem. Res.*, 2019, **52**, 3372–3383.
- 40 C. Hou, J. Wang, W. Du, J. Wang, Y. Du, C. Liu, J. Zhang, H. Hou, F. Dang, L. Zhao and Z. Guo, *J. Mater. Chem. A*, 2019, 7, 1346–13472.
- 41 X. Zang, C. Shen, Y. Chu, B. Li, M. Wei, J. Zhong, M. Sanghadasa and L. Lin, *Adv. Mater.*, 2018, **30**, e1800062.
- 42 R. Ge, J. Huo, M. Sun, M. Zhu, Y. Li, S. Chou and W. Li, *Small*, 2021, 17, e1903380.
- 43 R. Ma, W. Hao, X. Ma, Y. Tian and Y. Li, *Angew. Chem., Int. Ed.*, 2014, 53, 7310–7315.
- 44 Y. Ma, G. Guan, X. Hao, J. Cao and A. Abudula, *Renewable Sustainable Energy Rev.*, 2017, 75, 1101–1129.
- 45 A. Zunger, S.-H. Wei, L. Ferreira and J. E. Bernard, *Phys. Rev. Lett.*, 1990, **65**, 353.
- 46 P. Giannozzi, S. Baroni, N. Bonini, M. Calandra, R. Car, C. Cavazzoni, D. Ceresoli, G. L. Chiarotti, M. Cococcioni and I. Dabo, et al., J. Phys.: Condens. Matter, 2009, 21, 395502.
- 47 J. P. Perdew, K. Burke and M. Ernzerhof, *Phys. Rev. Lett.*, 1996, 77, 3865.
- 48 K. F. Garrity, J. W. Bennett, K. M. Rabe and D. Vanderbilt, Comput. Mater. Sci., 2014, 81, 446–452.
- 49 H. J. Monkhorst and J. D. Pack, Phys. Rev. B, 1976, 13, 5188.
- 50 L. Ward, R. Liu, A. Krishna, V. I. Hegde, A. Agrawal, A. Choudhary and C. Wolverton, *Phys. Rev. B*, 2017, 96, 024104.
- 51 Y. Jiang, D. Chen, X. Chen, T. Li, G.-W. Wei and F. Pan, *npj Comput. Mater.*, 2021, 7, 28.

PCCP

52 Z. Wu, J. Wang and H. Du, et al., Nat. Commun., 2023, 14, 2585.

- 53 X. Zhang, S. Sun and S. Wang, RSC Adv., 2020, 10, 43412-43419.
- 54 M. Kaupp, Angew. Chem., Int. Ed., 2001, 40, 3534-3565.
- 55 F. Viñes, C. Sousa, P. Liu, J. A. Rodriguez and F. Illas, J. Chem. Phys., 2005, 122, 174709.
- 56 C. R. Landis, T. Cleveland and T. K. Firman, J. Am. Chem. Soc., 1995, 117, 1859–1860.
- 57 M. Kaupp, J. Comput. Chem., 2007, 28, 320-325.
- 58 A. Nandy, D. B. Chu, D. R. Harper, C. Duan, N. Arunachalam, Y. Cytter and H. J. Kulik, *Phys. Chem. Chem. Phys.*, 2020, 22, 19326–19341.
- 59 W.-S. Ko, B. Grabowski and J. Neugebauer, *Phys. Rev. B:Condens. Matter Mater. Phys.*, 2015, **92**, 134107.
- 60 D. Farkas and A. Caro, J. Mater. Res., 2020, 35, 3031-3040.
- 61 B. Sharma, Y. S. Teh, B. Sadigh, S. Hamel, V. Bulatov and A. Samanta, *Comput. Mater. Sci.*, 2023, 230, 112486.
- 62 J. A. Moriarty, Theory and application of quantum-based interatomic potentials in metals and alloys/John A. Moriarty. [electronic resource], Oxford University Press, 2023.
- 63 M. Fey and J. E. Lenssen, *Fast Graph Representation Learning* with PyTorch Geometric, 2019.

- 64 A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai and S. Chintala, PyTorch: An Imperative Style, High-Performance Deep Learning Library, arXiv, 2019, preprint, arXiv:1912.01703, DOI: 10.48550/arXiv.1912.01703, https://arxiv.org/abs/1912.01703.
- 65 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, *Comput. Mater. Sci.*, 2013, 68, 314–319.
- 66 S. Pablo-García, S. Morandi, R. A. Vargas-Hernández, K. Jorner, Z. Ivković, N. López and A. Aspuru-Guzik, *Nat. Comput. Sci.*, 2023, 3, 433–442.
- 67 A. Hagberg, P. J. Swart and D. A. Schult, *Exploring network structure, dynamics, and function using NetworkX*, 2008, https://www.osti.gov/biblio/960616.
- 68 J. D. Hunter, Comput. Sci. Eng., 2007, 9, 90-95.
- 69 M. L. Waskom, J. Open Source Software, 2021, 6, 3021.
- 70 A. Stukowski, Modell. Simul. Mater. Sci. Eng., 2009, 18, 015012.
- 71 K. Momma and F. Izumi, *J. Appl. Crystallogr.*, 2011, 44, 1272–1276.