Supplementary Information (SI) for Physical Chemistry Chemical Physics.
This journal is © the Owner Societies 2025

Rapid Explainable GNN-Derived Structure-Property Relationships in Interstitial-Alloy Materials: a Demonstrated Workflow Using Mo2C

Supporting Information

Eduardo Aguilar Bejarano,[a,b,c] Luis Arrieta,[d] Mauricio Gutierrez,[e] Ender Özcan,[c] Simon Woodward, [a,b] Grazziela Figueredo*[f] and Ignacio Borge-Durán*[g]

[a] The GSK Carbon Neutral Laboratories for Sustainable Chemistry, The University of Nottingham, Jubilee Campus, Nottingham, NG7 2TU, UK.
[b] School of Chemistry, The University of Nottingham, University Park, Nottingham, NG7 2RD, UK.
[c] School of Computer Science, The University of Nottingham, Jubilee Campus, Nottingham, NG7 2TU, 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, UK.
[g] *Department of Physics, University of Alabama at Birmingham, Birmingham, AL, 35294, USA.*

*E-mail address:
jborgedu@uab.edu
g.figueredo@nottingham.ac.uk

CONTENTS

Table of Contents

Supporting Figures

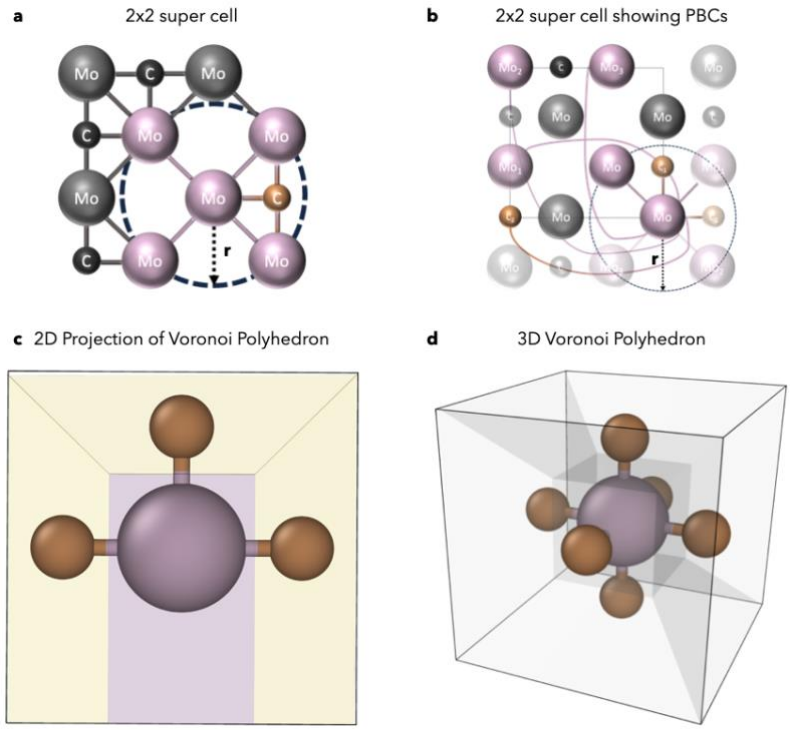

Figure S1. Scheme of graph representation used as input of the Crystal Graph Neural Network. **a)** Shows the radius threshold criteria used to determine whether or not two atoms share an edge, **b)** shows the radius threshold criteria applying periodic boundary conditions, **c)** shows the Voronoi polygons resulting from a Voronoi Tessellation algorithm applied to atoms that rely in 2D space, and **d)** shows the Voronoi polygons resulting from a Voronoi Tessellation algorithm applied to atoms that rely in 3D space.

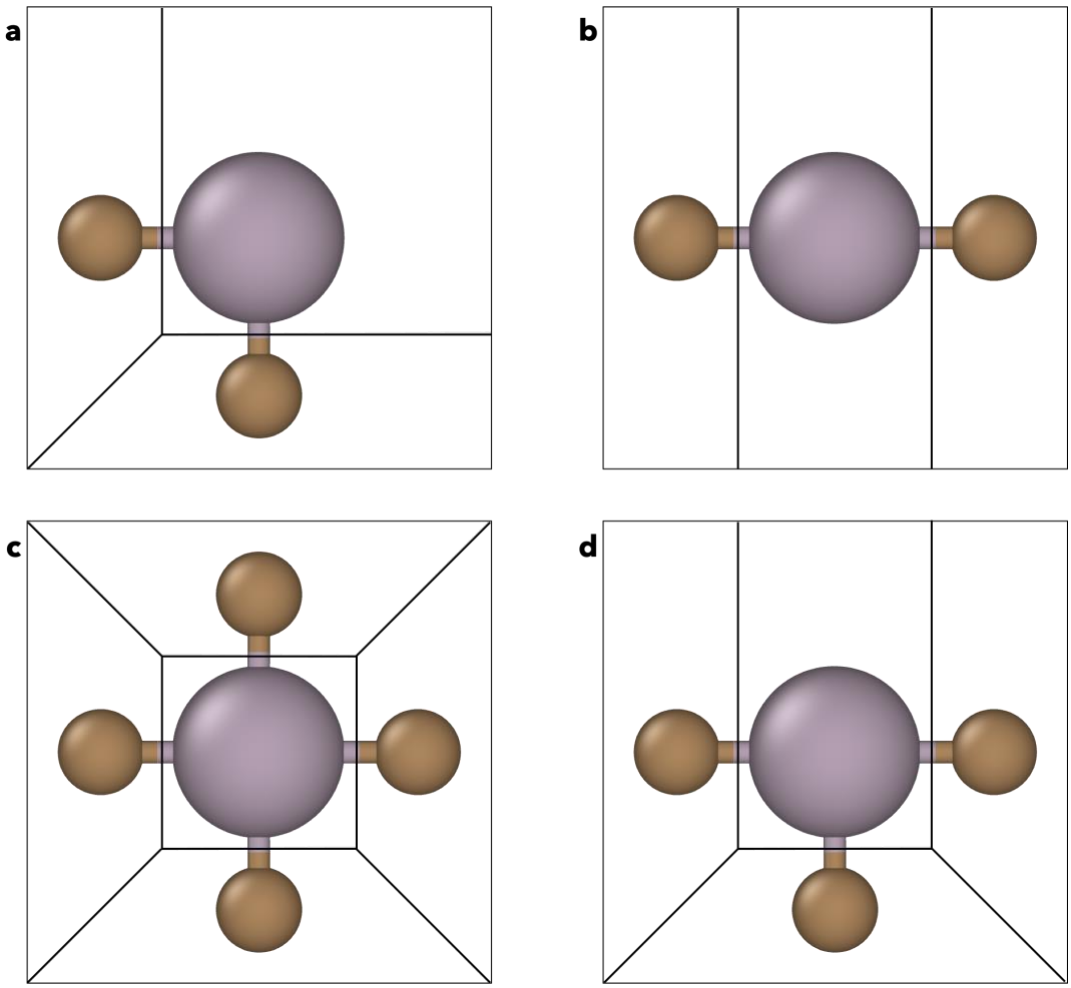

Fig. S2. Illustrative exemplars of Voronoi Polyhedrons for various atomic arrangements of carbide atoms around a molybdenum atom. The pictures show how the relative disposition of atoms makes the boundary of the central Mo polyhedron change.

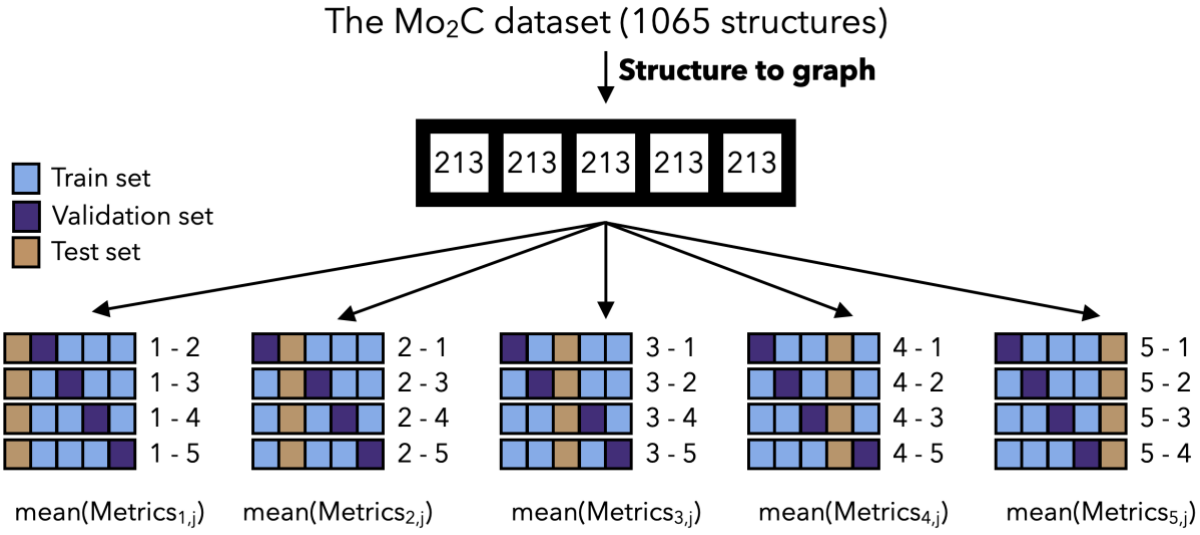

Fig. S3. Scheme of the inner cross validation approach for data splitting applied in this study.

Generating C/vacancy configurations in Mo_2C supercells for database creation

To diversify the atom arrangements, our approach started with a 1×1×1 cubic cell of MoC, using a 1:1 stoichiometry as a baseline to construct 2×2×2 supercells. We reduced half of the carbon atoms to achieve a 2:1 stoichiometry (2Mo:1C). We harnessed the method of constructing special quasi-random structures (SQS)[1] effectively mimicking the local atomic arrangements commonly seen in random solid solutions. This method is instrumental in depicting the inherent disorder characteristic of these materials accurately. The potential configuration count is determined as follows:

$$\binom{n}{k} = \frac{n!}{k!\,(n-k)k!}$$

Where:

- *n* represents the total number of possible positions for carbon atoms within the supercell or unit cell of the crystal structure.
- *k denotes the number of positions that are actually occupied by carbon atoms within the supercell or unit cell.*

Crystal Graph

The supercells were converted into graphs, where nodes represent the atoms and edges represent atom-to-atom interactions. To identify the type of atom, each node has a vector of dimensionality 2 containing the one-hot-encoded information of atom identity (Mo or C). For the edges, Voronoi Polyhedrons[2,3] were built around each atom within a cell. These 3D figures encapsulate the volume in space that is closer to the central point compared to any other point in the system. For a better understanding of this concept, Fig. S1 shows a Voronoi polygon for a series of atoms that lie in the 2D space and a Voronoi Polyhedron around a central point for a system of atoms in the 3D space.

The Voronoi Polyhedron is useful in deriving information on how the central point interacts with the other points in the system. We used atoms as central points, so that each atom within the supercell has an associated polyhedron. As shown in **Fig. S1**, each polyhedron is delimited by neighbouring polyhedrons, and such boundaries provide further information about the interaction between two points, or atoms in our case.[2] Particularly, we have quantified the interaction of atoms by calculating the solid angle. This parameter is a measurement of the amount of the field of view that is taken by an object being viewed from a specific point. This measure was used to add edge features to the crystal structures, as previously performed by Gu et al. to represent perovskites.[4]

By calculating the solid angle for a central atom and the polyhedron of a neighbouring atom, we can estimate the extent to which an atom is perturbed by the presence of that neighbouring atom, the topology of the crystal, and the relative position of atoms around a central atom. These factors are key to explaining the energetics of an atomic system (**Fig. S2**).[2–4] We expanded the information on the solid angle using a Gaussian Expansion method, which maps a variable *x* in R^1 to R^n, where *n* is a hyperparameter. These distributions follow Equation 4, where the difference between them is given by the μ value.

$$G(\Omega, \mu, \sigma) = exp\left(\frac{(\Omega-\mu)^2}{\sigma^2}\right)$$

We created the μ values from the interval [0, 3.5], with a step of 0.5; σ was set to 0.5 (because the series of numbers increases by this magnitude), and Ω is the solid angle. Due to the number of μ samples used, our expansion method takes Ω in R^1 and maps it to R^8. Therefore, each edge was given features of the expanded solid angle with a length of eight.

Model Training

Adam optimizer was used and the root mean squared error was adopted as the loss function. For each epoch, the loss of the validation set was supervised for two reasons: learning rate adjustment and early stopping. If no improvement in the validation loss were achieved for 7

subsequent epochs, the learning rate was multiplied by a factor of 0.7. The learning rate was initially set to 0.01, and the minimum learning rate allowed was 10^{-8}. If no improvement occurred in the validation set for 30 epochs, training was stopped. For the learning process, the training set was fed to the model in batches of 40 graphs, the loss and gradients were calculated for each batch, and weights and biases were updated after each batch. For each epoch, all training set was fed in batches, and an epoch was finalised once all batches had gone through the backpropagation process.

To assess robustness and generalisation of the models, we applied an inner cross-validation strategy (see **Fig. S3**). This splitting strategy consists in dividing the dataset into k folds and utilise each fold as test set, but also in each test set, each of the remaining k-1 folds is used as a validation set once, while the remaining k-2 folds are used as the training set. This way, for each test fold, a total of k-1 different training processes are performed, each with different training and validation sets. When all possible combinations of validation and training sets are achieved, then the test set is put back into the training set, and another set is taken out to be used as a test set. This process is repeated until all folds have been used as a test set. By doing this, it is possible to analyse the model generalisability when training on different datapoints. For our case, since we have created 5 folds, there would be a total of 5 test sets, while each test set will be evaluated after 4 different training processes using different validation and training sets, leading to a total of 20 different training processes.

Crystal Graph Explainer

The Crystal Graph Explainer is a novel tool to explain the relation between ensembles of atoms in a crystal structure and their contribution to the prediction of the property of the material by a Graph Neural Network. The nature of this tool goes further from former strategies, as they assumed no defects in the crystal structure and quantified the impact of having a certain atom type in a certain position. For the Crystal Graph Explainer, such assumption is not taken and therefore allows the quantification of the impact of having certain atom types in given relative arrangements to the overall properties of the structure.

The crystal graph explainer works by masking atoms from the graph structure before the pooling operation occurs. As pooling operation is useful to summarize the information contained on all the nodes, then this operation must take more information from those nodes that the GNN found to correlate the most with the target property. If these nodes are masked (hidden), then pooling cannot take information from such nodes, which leads to a different graph-level representation. As the graph representation is different, the prediction must be different. We define the attribution of the masked graph as the difference between the prediction of the complete graph and the masked graph. This way, the fragments that destabilize the crystal structure will have bigger attributions.

While this approach has been successfully applied to organic molecules, some issues may arise in the case of crystal strictures, including:

1. A fragment that needs to be explained may occur multiple times for the same central atom. For example, a molybdenum atom surrounded by 6 carbide atoms forming an octahedron is the super-position of three C-Mo-C fragments forming a 180° angle in the three different axes (x, y, and z) (**Fig. S4a**). When masking out all the matching fragments, the resulting structure would be equivalent to another structure. To fix that, the algorithm won't mask an atomic ensemble contained within a higher order ensemble. Therefore, only those C-Mo-C fragments forming a 180° angle with no other carbide atom around the central Mo will be masked.
2. A fragment that needs to be explained can occur many times within one single unit-cell in different sets of atoms (**Fig. S4b**). This will lead to an unfair comparison between attribution scores from different cells with different amounts of fragments removed. To fix this, we normalize the attribution value of the fragments by dividing it by the number of matching substructures that were removed from the cell.
3. A fragment that needs to be explained may happen more than once within the same unit cell and the matching fragments may overlap between them (**Fig. S4c**). In such case, the algorithm was coded so that it cannot remove twice the same atom. This way, if a central Mo originally matched the desired structure, but one of its carbons is removed because it

also matched the fragment but with another central Mo, then the first Mo will not be removed as the removed carbon causes such fragment to not be there anymore.

4. A fragment that needs to be explained does not occur at all in the unit-cell or supercell (**Fig. S4d**). In such cases, the algorithm does not take that observation into the analysis since it will generate an attribution score of 0 because of absence of the fragment and not because it is not important.

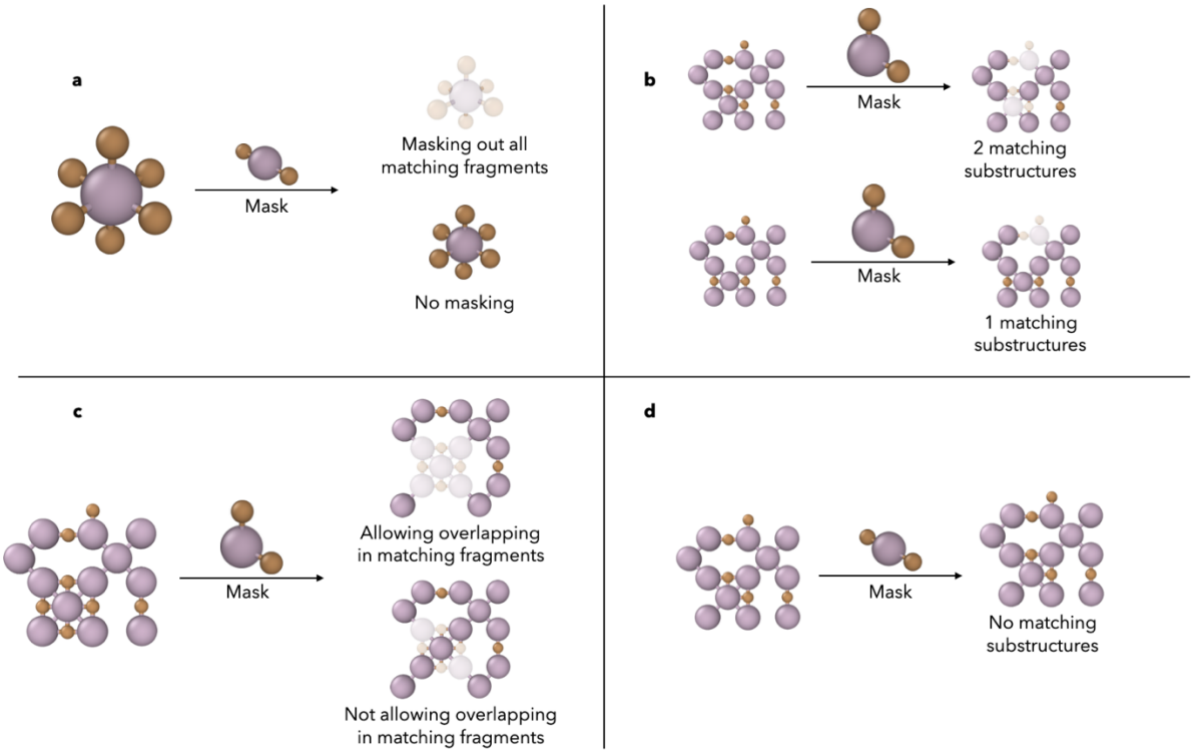

Fig. S4. Possible masking problems in crystal structure graphs.

Performance of the models on the 2x2x2 cells

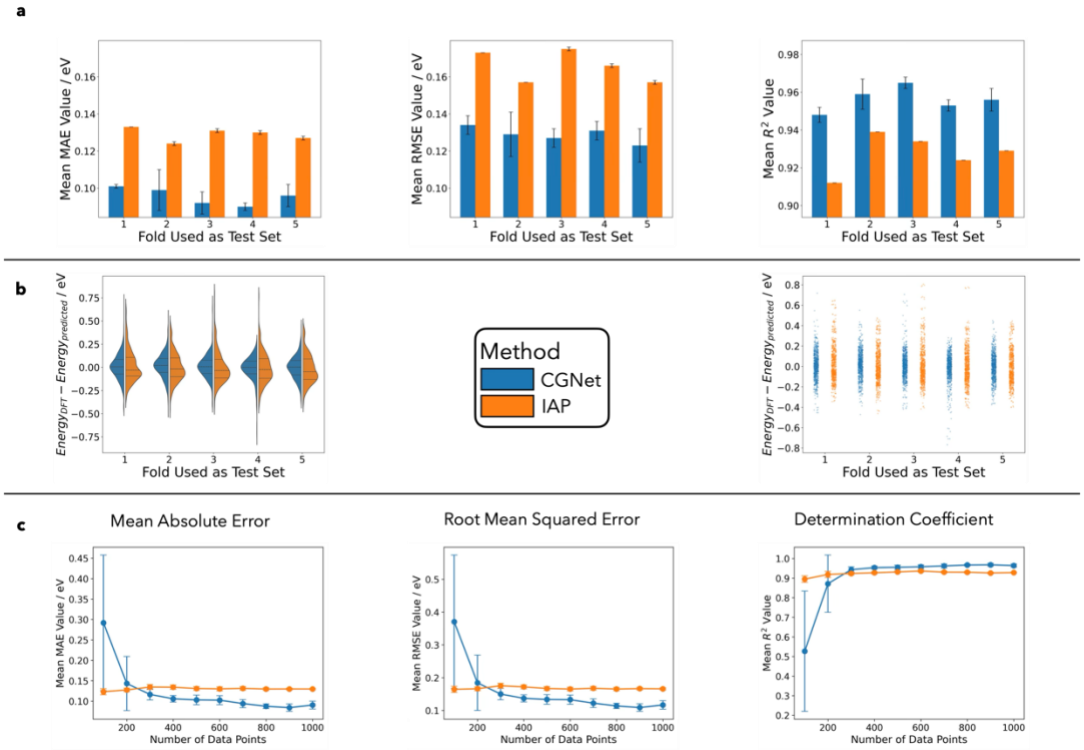

Fig. S5. Summary of results obtained by Crystal Graph Neural Network and Interatomic Potential approaches on the Mo_2C dataset. The reported values are shown for each test fold separately, using the mean of the metrics for the value in the bar and the standard deviation as error bars. a) Shows the mean metrics obtained for each test fold and standard deviation as error bars. b) Shows the error distribution of each test set. C) Shows the learning curve of the methods for different metrics.

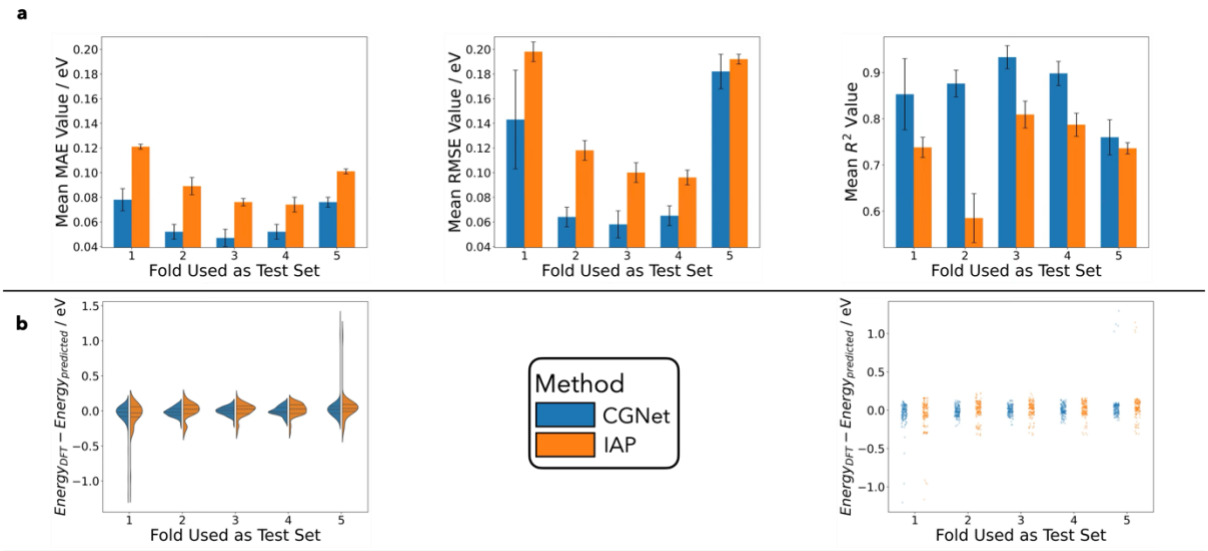

Fig. S6. Summary of results obtained by Crystal Graph Neural Network and Interatomic Potential approaches on the Ti_2C dataset. The reported values are shown for each test fold separately, using the mean of the metrics for the value in the bar and the standard deviation as error bars. a) Shows the mean metrics obtained for each test fold and standard deviation as error bars. b) Shows the error distribution of each test set.

References

(1) Zunger, A.; Wei, S.-H.; Ferreira, L. G.; Bernard, J. E. Special Quasirandom Structures. *Phys. Rev. Lett.* **1990**, *65* (3), 353–356. https://doi.org/10.1103/PhysRevLett.65.353.

(2) Ward, L.; Liu, R.; Krishna, A.; Hegde, V. I.; Agrawal, A.; Choudhary, A.; Wolverton, C. Including Crystal Structure Attributes in Machine Learning Models of Formation Energies via Voronoi Tessellations. *Phys. Rev. B* **2017**, *96* (2), 024104. https://doi.org/10.1103/PhysRevB.96.024104.

(3) Jiang, Y.; Chen, D.; Chen, X.; Li, T.; Wei, G.-W.; Pan, F. Topological Representations of Crystalline Compounds for the Machine-Learning Prediction of Materials Properties. *npj Computational Materials* **2021**, *7* (1), 28. https://doi.org/10.1038/s41524-021-00493-w.

(4) Gu, G. H.; Jang, J.; Noh, J.; Walsh, A.; Jung, Y. Perovskite Synthesizability Using Graph Neural Networks. *npj Computational Materials* **2022**, *8* (1), 71. https://doi.org/10.1038/s41524-022-00757-z.