

Cite this: *J. Mater. Chem. A*, 2025, **13**, 33385

Charting the chemical space of Zintl phases with graph neural networks and bonding insights

Rinkumoni Chaliha,^{†ab} Manish Kothakonda,^{†b} Cheng-Wei Lee,^b Jeffrey N. Law,^c Qian Yang,^d Svilen Bobev,^e and Prashun Gorai^{*ab}

A large number of Zintl phases have been discovered by solid-state chemists driven by empirical knowledge, chemical intuition and in some cases, through serendipitous accidents. These discoveries have only scratched the surface, given the vast compositional and structural diversity that Zintl phases can accommodate. The large chemical space of Zintl phases, as well as intermetallic compounds in general, remain under-explored. Here, we use graph neural networks and the upper bound energy minimization approach to efficiently scan a large chemical space of >90 000 hypothetical Zintl phases and accurately discover 1810 new thermodynamically stable phases with 90% precision, as validated with first-principles calculations. We show that our approach is more than 2× more accurate in predicting DFT stability than M3GNet (40% precision) on the same dataset. Using a random forest model and SHAP analysis, we demonstrate the critical role of ionic bonding in the thermodynamic stability of Zintl phases. Our results not only expand the known chemical landscape of Zintl phases but also highlight the efficacy of machine learning frameworks combined with domain knowledge in uncovering chemically meaningful insights across complex intermetallics.

Received 31st July 2025
Accepted 29th August 2025

DOI: 10.1039/d5ta06210a

rsc.li/materials-a

1 Introduction

Much of the vast chemical space of plausible inorganic compounds remains unexplored. The need to systematically explore this chemical space is shared by experimental and computational researchers. Currently, the number of entries in the Inorganic Crystal Structure Database (ICSD) is close to 300 000, which represents only a small part of the inorganic chemical space. An estimated more than >10¹² valence-balanced compounds are plausible, considering up to quaternary compositions.¹ There is an even larger chemical space of inorganic disordered phases and alloys. However, not all plausible compounds are thermodynamically stable or synthetically accessible. By charting even subsets of the inorganic chemical space, we not only discover new materials, but also develop structure–property relationships that will enable systematic design of functional materials. In this study, we present a machine learning (ML)-driven exploration of the Zintl chemical space, leveraging a graph neural network framework to discover new stable Zintl compounds. We not only expand the known library of stable inorganic compounds but also provide

a blueprint for how data-driven approaches can uncover chemically meaningful trends and structure–property relationships for materials design.

Zintl phases are a fascinating class of intermetallics, typically composed of more electropositive alkali, alkaline earth, or rare earth metals combined with more electronegative p-block elements.^{2,3} The electrons donated by the more electropositive elements are used to form covalent and polar covalent bonds among the more electronegative elements in a way that satisfies electron counting within the Zintl–Klemm formalism.² Consequently, the structures of Zintl phases are often characterized by “spectator” cations that form ionic bonds with polyanionic frameworks composed of the more electronegative elements which are bound *via* covalent, polar covalent, and sometimes, metallic bonds. The combination of complex structures and unique bonding gives rise to a broad spectrum of electronic, thermal, and mechanical properties.^{4–8} Despite their promising attributes, the discovery of new Zintl phases has traditionally relied on empirical rules, chemical intuition, and exploratory solid-state synthesis – approaches that, while fruitful, are inherently limited in scope and scale.

The structural and compositional complexity of Zintl phases arises from their unique bonding schemes, which give rise to an array of stable and metastable structures.^{2,9,10} However, this same complexity makes it challenging to predict their stability and structure–property relationships. First-principles calculations based on density functional theory (DFT) are relatively accurate, but computationally expensive and thus impractical

^aRensselaer Polytechnic Institute, Troy, New York, USA. E-mail: gorai@rpi.edu^bColorado School of Mines, Golden, CO, USA^cNational Renewable Energy Laboratory, Golden, CO, USA^dUniversity of Connecticut, Storrs, CT, USA^eUniversity of Delaware, Newark, DE, USA[†] Equal contribution.

for exhaustive screening across large chemical spaces. Consequently, the vast majority of chemically plausible Zintl compounds remain unexplored, representing a largely untapped reservoir of potentially novel materials. Discovery of new Zintl phases using DFT high-throughput screening¹¹ and ML approaches¹² have been attempted.

Recent advances in ML, and in particular the development of graph neural networks (GNNs), offer a powerful alternative to traditional methods by enabling the prediction of thermodynamic stability and material properties directly from crystal structures or chemical compositions with significantly reduced computational cost. GNNs are especially well-suited to this task due to their ability to naturally encode atomic connectivity and local chemical environments. Modern GNNs trained on data from materials databases predict the thermodynamic (convex hull) stability with high accuracy, with typical errors lower than the “chemical accuracy” of 1 kcal mol^{−1} (43 meV per atom).^{13–17} The prediction accuracy, when extrapolating to hypothetical structures, is improved by incorporating high-energy structures in the training data in addition to ground-state and near ground-state structures.¹⁸

The high accuracy of GNN models in predicting thermodynamic stability requires optimized (relaxed) crystal structures as inputs. Typically, DFT is used to optimize input structures, which is computationally expensive, especially for hypotheticals and a large number of input structures. The success of using an unrelaxed structure as input depends on how close the structure is to the relaxed version on the potential energy surface.¹⁹ Scale-invariant GNN models are only tolerant to changes in cell volume, but not to cell shape and atomic positions.²⁰ Machine-learned interatomic potentials (MLIPs) are becoming increasingly sophisticated with almost DFT-level accuracy, but limitations of the training data – both in chemical and structural diversity, and far-from-equilibrium sampling – restrict their “universal” accuracy in performing structural relaxations. Law *et al.* proposed a novel approach to discover thermodynamically stable phases by using a scale-invariant GNN model that predicts the volume-only relaxed (constrained structure) energy from input unrelaxed structures, avoiding the challenges of predicting fully relaxed energy.²¹ The approach is named Upper Bound Energy Minimization (UBEM) as the volume-relaxed energy, by design, exceeds the fully relaxed energy, yielding an upper bound. Consequently, if the volume-relaxed structure is thermodynamically stable (on the convex hull), the fully relaxed structure will assuredly be stable.

In this work, we present a ML-driven exploration of the Zintl chemical space, leveraging a GNN framework to screen over 90 000 hypothetical Zintl phases. Using the UBEM strategy, we identify 1810 new thermodynamically stable Zintl phases with a validation precision of 90% when benchmarked against DFT. Our GNN model not only outperforms existing MLIPs such as M3GNet,²² which achieves only 40% precision on the same dataset, but also demonstrate generalizability across different chemistries and complex structural frameworks within the Zintl chemical space. To gain deeper composition–structure–property insights, we further employ a random forest model combined with SHAP (SHapley Additive exPlanations) analysis

and find that the strength of ionic bonding plays a dominant role in determining the thermodynamic stability of Zintl phases, reaffirming our chemical intuition within a quantitatively rigorous framework.

2 Results and discussion

2.1 Prediction of stable Zintl phases with UBEM approach

Our data-driven search for stable Zintl phases began with a curated dataset of 824 pnictide-based Zintl phases extracted from the Inorganic Crystal Structure Database (ICSD).²³ The identification of Zintl phases, which relies only on their compositional make up, is described in the Computational methods (Section 3.1.1). This dataset of structural prototypes encompasses binary, ternary, and quaternary compositions, representing a broad range of structural frameworks. We systematically explored chemical substitutions by decorating these parent structures with elements from Groups 1, 2, 12, 13, 14, and 15, including only Mn among the transition metals (encountered most commonly as Mn²⁺) for isovalent substitutions (Section 3.1.3). This strategy yielded a large chemical space comprising over 90 000 candidate hypothetical structures.

To efficiently identify thermodynamically stable phases within this large search space, we employed the UBEM approach.²¹ This approach leverages a scale-invariant GNN model to predict DFT volume-relaxed energies using unrelaxed crystal structures as input. The scale-invariant architecture normalizes input structure volumes, rendering the model less sensitive to volume distortions^{20,21} and capable of accurately predicting volume-relaxed energies even for unrelaxed geometries. Details of the GNN model, training, and phase stability analysis are provided in the Computational methods (Section 3.2).

To validate the foundational assumption of UBEM—that the volume-relaxed energy provides an upper bound to the fully-relaxed DFT energy—we selected a representative subset of structures that preserved the compositional and structural diversity of the full dataset. As shown in Fig. 1a, we plotted the DFT volume-relaxed energy against the fully-relaxed energy for this subset. The results confirm that the volume-relaxed energies are consistently higher, thus supporting their use as upper bounds. Notably, the plot also highlights cases where the volume-relaxed and fully-relaxed energies are nearly identical, indicating minimal changes in cell shape and ionic positions during relaxation.

By including both types of structures that exhibit large and minimal energy differences (Fig. 1a) during GNN training, we ensured the model could generalize across a range of structural relaxation behaviors.¹⁸ Fig. 1b shows the learning curve of the GNN surrogate model as a function of training dataset size. For each training split, 10% of the data was reserved for testing, and model performance was evaluated using mean absolute error (MAE). The final model, trained on 6571 volume-relaxed structures, achieved a low test MAE of 27 meV per atom. Additionally, a parity plot comparing model predictions to DFT-calculated volume-relaxed energies (Fig. 1c) reveals excellent agreement, demonstrating the surrogate model's accuracy. These results






Fig. 1 (a) Comparison of DFT fully-relaxed and volume-relaxed energies of 2584 hypothetical Zintl structures. (b) Learning curve of scale-invariant GNN model trained on DFT volume-relaxed energies. The standard deviation (shown as error bars) is calculated from 4 different models with non-overlapping test sets. (c) GNN-predicted volume-relaxed energy is in good agreement with DFT volume-relaxed energies with a low mean absolute error of 27 meV per atom.

confirm that training on volume-relaxed structures alone is sufficient, offering a more computationally efficient alternative to prior methods that relied on fully-relaxed DFT structures.

We applied the trained GNN model to predict the volume-relaxed energy of all >90 000 chemically decorated structures. For each composition, we identified the candidate with the lowest predicted energy as the representative upper bound minimum structure. These minima were then subjected to thermodynamic stability analysis (Section 3.2.6) by computing decomposition energies (E_{decomp}) relative to competing phases from the ICSD, using total energies available in the NREL Materials Database.²⁴

While $E_{\text{decomp}} \leq 0$ indicate stability, the use of volume-relaxed energies (which are upper bounds) can lead to false negatives, *i.e.*, cases where a structure appears metastable but could become stable upon full relaxation. To mitigate this risk, we adopted a conservative threshold, selecting structures with E_{decomp} below 10 meV per atom as potentially stable. This criterion yielded 1810 promising new Zintl phases that are likely stable against decomposition into competing phases and warrant further investigation.

2.2 Prediction of stable Zintl phases with M3GNet

To benchmark the performance of the UBEM method against other established techniques, we employed the M3GNet interatomic potential (IAP) developed by Chen *et al.*²² This is a GNN-based IAP and is trained on the extensive Materials Project database,²⁵ which includes energies, forces, and stresses from DFT structural relaxations. Using pre-trained M3GNet, we performed full structural relaxations for all >90 000 decorated structures in our search space. Fig. S1 presents a comparison between total energies predicted by M3GNet and those obtained from DFT calculations. While the data shows a moderate correlation, M3GNet tends to systematically underestimate the total energies relative to DFT. This systematic underestimation

poses a risk of false positives in stability predictions, as some structures may appear more stable than they actually are.

Following relaxation with M3GNet, we selected the lowest-energy structure for each composition. Using these minima, we constructed convex hulls by including competing phases from the Materials Project database. Unlike UBEM, which relies on volume-relaxed energies and therefore requires a E_{decomp} threshold, M3GNet directly predicts fully-relaxed total energies. As such, any structure with a negative E_{decomp} is considered stable without applying an additional stability tolerance. Using this criterion, the M3GNet IAP predicts 1754 stable Zintl phases within the same search space explored by UBEM. In addition, M3GNet identifies 242 extra stable phases containing Yb, which were not included in the UBEM search. A complete list of these phases is provided in the GitHub Repository (see Data availability).

2.3 DFT validation of thermodynamic stability predictions

To evaluate the robustness of thermodynamic stability predictions made by the UBEM and M3GNet approaches, we performed full DFT structural relaxations (without constraints) on all 1810 UBEM-predicted and 1754 M3GNet-predicted stable structures. For UBEM candidates, we used the same DFT computational setup as the NREL Materials Database.²⁴ For M3GNet candidates, we adopted the DFT settings consistent with the Materials Project.²⁵ Following full relaxations, we conducted self-consistent convex hull analyses (see Computational methods) for each set of predicted structures. For UBEM-predicted phases, we constructed convex hulls using fully relaxed energies and competing phases from the NREL Materials Database. Similarly, for M3GNet-predicted phases, we used relaxed energies and competing phases drawn from the Materials Project database. From these analyses, we calculated E_{decomp} of all fully-relaxed structures.

To assess the predictive performance of each method, we compared E_{decomp} from machine-learned predictions to those






Fig. 2 Comparison of DFT-computed (fully-relaxed structures) decomposition energy (E_{decomp}) with (a) UBEM and (b) M3GNet predictions of E_{decomp} . The data distributions are shown as histograms.

from DFT, as shown in the parity plots in Fig. 2. We assess the predictive power in terms of precision, calculated as the fraction of predicted stable phases that remain stable after full DFT relaxation, *i.e.*, those with negative E_{decomp} .

As shown in Fig. 2a, 1634 out of 1810 UBEM-predicted phases (with initial $E_{\text{decomp}} < 10$ meV per atom) were confirmed to be thermodynamically stable after full DFT relaxation. This corresponds to a high precision of 0.90, demonstrating the effectiveness of using volume-relaxed energies as upper bounds for screening stable materials. This result highlights UBEM's capability to efficiently and accurately predict stability. Moreover, precision could be further improved by lowering E_{decomp} threshold (*e.g.*, below 10 meV per atom), which would reduce false positives at the expense of reduced number of candidate materials.

In contrast, only 668 of the 1754 structures predicted as stable by M3GNet were validated as stable by DFT (Fig. 2b), resulting in a much lower precision of 0.38. This reduced precision is consistent with the earlier observation that M3GNet tends to underestimate total energies, increasing the likelihood of false positives in stability predictions.¹⁶

2.4 Experimental validation of stability predictions

As DFT validation is inherently limited due to the functionals used, the reliability of UBEM stability predictions is further highlighted by independent experimental confirmations. Although our GNN model was trained exclusively on the 2016 version of the ICSD, it successfully predicted the stability, not only of the compositions but also of the correct crystal structures, of several Zintl phases that were not present in the training dataset but have since been synthesized and reported in the experimental literature over the past several years. This highlights the predictive power of our approach, extending beyond retrospective fitting to forward discovery. Notably, several Zintl phases, including $\text{Ca}_{14}\text{AlBi}_{11}$ ($I4_1/acd$),²⁶ $\text{Ca}_{14}\text{-MgBi}_{11}$ ($I4_1/acd$),²⁷ CaIn_2As_2 ($P6_3/mmc$),²⁸ BaIn_2As_2 ($P2/m$),²⁸ Ca_3InAs_3 ($Pnma$),²⁹ Sr_3InAs_3 ($Pnma$),³⁰ RbMn_4As_3 ($P4/mmm$),³¹

and Ba_2ZnP_2 ($Ibam$),³² have been experimentally synthesized in recent studies, and all were accurately predicted as stable by our model. These successes serve as compelling evidence of the model's reliability. Interestingly, all of these phases were identified by the UBEM approach, with Sr_3InAs_3 , Ca_3InAs_3 , and Ba_2ZnP_2 also being predicted by M3GNet. This overlap reinforces the broader value of both models but also illustrates the higher precision of the UBEM approach in practice. However, there are instances where our model correctly predicts a stable composition but not its ground-state structure, *e.g.*, SrIn_2As_2 ,²⁸ our model identifies the $R\bar{3}2/m$ structure as the ground state, whereas the experimentally observe $P6_3/mmc$ phase was predicted to be 0.034 eV per atom higher than the ground-state. While this reflects a slight energy overestimation, it highlights the ability of the model to capture energetically competitive polymorphs. Additionally, our model also exhibits false negatives. For example, KCaBi ($P4/nmm$)³³ and KMgBi ($P4/nmm$)³⁴ are predicted to be 0.021 eV per atom and 0.091 eV per atom above the convex hull, respectively, despite their experimental synthesis. These false negatives occur due to the use of volume-relaxed energies to predict stability. Importantly, subsequent full DFT relaxations confirmed that both structures as thermodynamically stable, aligning with experimental findings. Despite these limitations, the model maintains a very low false-positive rate, underscoring its reliability for high-throughput screening of new Zintl phases.

Fig. 3a compares the elemental distributions of newly predicted stable Zintl phases with those in the ICSD, providing a global view of the discovery landscape accessed through our decoration strategy. The results reveal a predominance of heavy elements such as Rb, Cs, Hg, Tl, Pb, and Bi in the predicted structures. The ratio R of the number of newly discovered phases containing a specific element to the number of known ICSD phases containing the same element is the highest for Tl, which possibly reflects a historical bias in the experimental discovery efforts due to acute toxicity of Tl. Similarly, experimental efforts to discover Rb- and Cs-containing Zintl phases could have been impeded by the highly air-sensitive nature of






Fig. 3 (a) Element-wise heatmap of the ratio (R) of the number of newly-predicted Zintl phases containing a specific element to the number of ICSD phases containing the same element. (b) UpSet plot showing intersecting sets of UBEM and M3GNet predicted and DFT-confirmed number of stable phases. Filled dots in each row indicate which sets are intersecting, and the corresponding bar height shows the number of Zintl phases in that intersection, analogous to how Venn diagrams represent overlapping subsets. The horizontal bars are the number of Zintl phases predicted or confirmed with each method.

those compounds. The discovery of Zintl phases with heavier elements is particularly notable for thermoelectrics, as heavy elements are known to suppress lattice thermal conductivity.³⁵ In Fig. S2, we show the distribution of the 30 most common space groups among our UBEM-predicted and DFT-confirmed structures. 100 of these structures crystallize in the $P6_3/mmc$ (No. 194) space group. Numerous 1-1-1 Zintl phases in $P6_3/mmc$ have previously been predicted as topological insulators or Dirac semimetals.³⁶ We also discovered 35 structures that crystallize in $I4_1/acd$ space group (No. 142). Notably, 14-1-11 Zintl phases that crystallize in this tetragonal space group are among the most efficient high-temperature thermoelectrics, with figure-of-merit zT values approaching or even exceeding unity.^{37–39} Thus, beyond confirming stability, our predictions may guide the development of functional materials with desirable transport properties.

To further illustrate the comparative performance of UBEM and M3GNet, we present an UpSet plot in Fig. 3b. The plot reveals that 407 phases were predicted by both UBEM and M3GNet, of which 260 were confirmed to be stable *via* DFT. This agreement likely reflects the shared GNN-based framework underlying both models, yet UBEM consistently shows higher

precision, especially in cases independently validated by experiment.

In summary, the alignment of our predictions with recently synthesized compounds, despite their absence from the training data, offers strong real-world validation of our approach. It demonstrates that machine learning models, when carefully trained and systematically deployed, can serve as powerful tools for accelerated materials discovery.

2.5 Interpretable ML to understand Zintl phase stability

While neural network models like UBEM and M3GNet are highly effective at predicting thermodynamic stability, they often operate as black boxes, which impedes interpretation due to the complexity of learned representations.⁴⁰ However, for materials discovery to be not only fast but also rational, it is crucial to understand why a material is predicted to be stable. In particular, interpreting stability in terms of fundamental chemical features can guide more targeted materials exploration, optimization, and design.

In Zintl phases the interplay of ionic (electrostatic interactions) and covalent bonding critically shapes both their crystal structures and functional properties like thermoelectric performance. To understand this interplay, we develop an interpretable Random Forest (RF) regression model,⁴¹ trained on a chemically diverse and expanded dataset that includes both newly discovered Zintl phases and those from the ICSD. The goal was to predict formation energy (E_f), a robust and transferable metric of material stability, directly from chemistry-informed features.

E_f is selected over E_{decomp} because it is independent of the stability of other (competing) phases and can be inferred purely from its elemental and structural attributes, making it broadly applicable across compositions and structure types.¹⁴ Stability in Zintl phases is thought to arise from a balance of two main contributions: (i) ionic bonding, which are essentially electrostatic interactions, between cationic and anionic sublattices, (ii) covalent and polar covalent bonding within the anionic framework and between cation–anion sublattice.³ Accordingly, we engineer features to represent both contributions, creating a set of 26 interpretable descriptors (detailed in the SI).

These features include ionic bonding descriptors such as cation charge density, cation charge-to-size ratio, and electronegativity differences between cations and anions, as well as the sum of ionic radii. Covalent bonding is captured through features such as covalent bond density, weighted by electronegativity and orbital energy, as a proxy for bond strength. Inspired by the classical van Arkel–Ketelaar triangle, we included average electronegativity and its standard deviation to differentiate between metallic, ionic, and covalent bonding characters. Other complementary descriptors included average valence electron count, atomic mass, atomic density, and CrystalNN-based local environment fingerprints.⁴²

The RF model, trained on UBEM- and M3GNet-predicted stable phases, and ICSD phases, achieves a 5-fold cross-validated mean absolute error (MAE) of 52 meV per atom, showing good agreement with DFT-calculated formation



energies (Fig. S3). Feature importance analysis revealed that several of the 26 features had negligible impact, as indicated by near-zero SHAP (SHapley Additive exPlanations) values (Fig. S4). We applied an iterative backward-elimination approach, removing the least important features based on updated SHAP values, while simultaneously monitoring MAE and tuning model hyperparameters.

As shown in Fig. 4a, the lowest MAE (49 meV per atom) is obtained with 12 features. However, a nearly equivalent MAE (50 meV per atom) is achievable with just 10 features, offering a favorable trade-off between performance and model simplicity. The 10-feature model used 200 decision trees, each with a maximum depth of 15, and considered 6 features per split. In contrast, the 12-feature model required deeper trees (depth of 20) and greater complexity, with minimal accuracy gain.

To evaluate generalizability, we examine the learning curves for the 10-feature model (Fig. 4b). The steadily decreasing MAE with increasing training data and the small gap between training and test errors confirm generalization without overfitting. A low final training error (20 meV per atom) at large data sizes indicates the absence of underfitting.⁴³ Although the learning curves in Fig. 4b show that errors will continue to decrease with more training samples, the MAE is low enough (50 meV per atom) for reliable feature importance interpretation, which is the primary aim of our analysis.

A parity plot comparing RF-predicted and DFT-calculated formation energies (Fig. 4c) shows good agreement, with a few outliers. The top 10 largest outliers (highlighted in red) predominantly involve Mn-containing compounds (8 out of 10). These discrepancies likely arise from the lack of explicit magnetic descriptors in the model. Another notable outlier, LiAs₇, showed a residual of 0.55 eV per atom, much higher than its heavier analogs (NaAs₇ ~ 0.37 eV per atom, KAs₇ ~ 0.11 eV per atom, RbAs₇ ~ 0.11 eV per atom), likely due to weak Li-As covalent interactions (bond length 2.79 Å) not captured by the model. Similarly, the higher energy zinc blende structure of RbSb (*F43m*) was significantly mispredicted (residual 0.44 eV per atom). This discrepancy may be due to the presence of excess electrons in zinc blende RbSb, where Sb has a -3 charge state. Interestingly, the model performs better in predicting E_f of the ground-state β -RbSb (*P2₁/c*, LiAs structure type), which is electron-balanced, resulting in a lower residual of 0.25 eV per atom.

To further probe model interpretability, we analyze SHAP values to determine the contribution of each feature to the predicted E_f (Fig. 5). SHAP summary plots rank features by importance and indicate whether high or low feature values increase stability, *i.e.*, lower E_f .⁴⁴ The colorbar in Fig. 5 indicates normalized magnitude of individual features ranging from blue (low) to red (high). Since E_f is more negative for more stable phases, a negative SHAP value reflects a stabilizing contribution. The top features are dominated by descriptors of ionic bonding: cation charge density per atom, average cation charge-to-size ratio, difference between maximum and minimum electronegativity, and average electronegativity difference between cations. These features show strong influence of




Fig. 4 (a) 5-Fold cross-validated MAE as a function of the number of features in the RF model. (b) Learning curve of the RF model with 10 features. (c) Parity plot comparing RF-predicted formation energy (E_f) with DFT-calculated E_f for 2444 Zintl phases. 5-Fold cross-validated MAE is 50 meV per atom.

electrostatic interactions in stabilizing Zintl phases. This is further supported by the negative direction of impact of the sum of ionic radii of elements with maximum and minimum electronegativity on E_f , as electrostatic interactions decrease with distance between ions. Increased phase stability with both low atomic density and atomic mass are indicators of stronger ionic






Fig. 5 SHAP values of the 10 features of the RF model (Fig. 4). The colormap denotes the feature values (low to high). Ionic and covalent bonding features are marked by “i” and “c” in front of the feature labels. EN is electronegativity and BD is bond density.

bonding favoring stability. Lower atomic density lowers repulsive electrostatic interactions between ions. Interestingly, only one covalent bonding descriptor – covalent bond density weighted by electronegativity – ranks among the top 10. This feature negatively impacts stability at high values, especially in light-element compounds, suggesting that excessive covalent bonding may sometimes destabilize Zintl phases.

To understand the interplay between ionic and covalent bonding in stabilizing Zintl phases, we create a classical van Arkel–Ketelaar bonding map (Fig. 6). The dashed lines are simply a guide to the eye. We find that points with high average electronegativity and low standard deviation in




Fig. 6 A van Arkel–Ketelaar bonding map (average electronegativity vs. standard deviation in electronegativity) of stable Zintl phases. Increased ionic character leads to more stable phases, as indicated by the more negative SHAP values of average electronegativity (colormap).

electronegativity—indicative of strong covalent bonding character – tend to make the Zintl phases less stable, as suggested by the predominance of positive SHAP values (of electronegativity) in this region. Conversely, the region with moderate average electronegativity values and high standard deviation, which is indicative of ionic bonding character, positively impacts stability (negative SHAP values). These insights are reinforced by interaction plots (Fig. S5a and b) showing that while moderate covalent bonding can coexist with ionic interactions, excessive covalency may reduce phase stability, particularly in compounds with lighter cations.

Our RF model provides not only accurate predictions of formation energy but also interpretable insights into the underlying chemistry of Zintl phases. Through systematic feature engineering and SHAP analysis, we demonstrate that ionic bonding plays a more decisive role than covalent interactions in determining the stability of Zintl phases. This work highlights how interpretable machine learning and domain knowledge can move us beyond “black-box” predictions to provide chemically meaningful guidance for future materials discovery.

3 Computational methods

A schematic of the computational workflow for identifying stable Zintl phases using the UBEM approach is shown in Fig. S6.

3.1 Prototype library and chemical decoration

To systematically explore the vast chemical space of Zintl phases, we performed ionic substitution (chemical decoration) using known ICSD Zintl phases as prototype structures to create hypothetical structures.^{11,45} Below, we describe the methodology used to curate the prototype library and the decoration rules applied to generate hypothetical compounds. Each hypothetical structure retains the connectivity and symmetry of its prototype, while differing in elemental makeup.

3.1.1 Identification of Zintl pnictide prototypes. We focused exclusively on Zintl pnictides, defined here as compounds containing one or more of the group 15 elements P, As, Sb, or Bi as anions. To identify Zintl phases from the Inorganic Crystal Structure Database (ICSD),²³ we applied a set of empirical chemical and structural criteria that allows selection of “Zintl” phases without considering their electronic structures, which is a suitable approach for large-scale explorations such as the one presented in this study. We acknowledge that classification of a compound strictly as a Zintl phase requires detailed structural analysis and electron counting, and even then, the classification is subject to scrutiny. After all, Zintl phases lie at the border of metallic intermetallics and salt-like compounds, which lends itself to sometimes unclear electron counting rules. Our goal here is to broadly identify compounds that may be classified as Zintl phases,^{3,9,46,47} ultimately, the classification does not affect the approach or general conclusions of this work.



For cations, we considered a specific subset of elements, grouped based on their electronegativity and typical roles in Zintl chemistry: (1) high-electropositive cations: Li, Na, K, Rb, Cs, Be, Mg, Ca, Sr, Ba, Eu, Yb, and (2) less-electropositive cations: Mn, Zn, Cd, Hg, Al, Ga, In, Tl, Si, Ge, Sn, Pb, and Be. Be was included in both groups due to its dual role as a spectator cation or as part of covalently bonded anionic framework, depending on the phase. We retained compounds that contain at least one of the above cations and one of the pnictogens (P, As, Sb, Bi). To exclude purely ionic materials, we enforced additional structural filters based on bonding: for binary Zintl pnictides, we selected ordered compounds that: (1) contain at least one high-electropositive cation, and (2) exhibit at least one pnictogen–pnictogen bond in the first nearest neighbor shell, which we define as all atoms within the shortest distance to a neighboring atom (d) and up to $1.15d$.⁴⁸ For ternary and quaternary compounds, we excluded structures lacking high-electropositive cations unless anion–anion bonding was explicitly observed using the same neighbor search method. Following these criteria, we identified 824 Zintl ordered pnictides in the 2016 version of the ICSD.

3.1.2 Prototype library. To create prototype structures for decoration, we grouped these 824 Zintl compounds based on a chemistry-agnostic classification using nominal oxidation state-aware composition type and space group. Each element in the structure was mapped to a generic placeholder based on its nominal oxidation state: a = P, As, Sb, Bi (pnictogens), b = Li, Na, K, Cs, Rb, c = Be, Mg, Ca, Sr, Ba, Zn, Cd, Hg, Mn, Yb, d = Al, Ga, In, Tl, and e = Si, Ge, Sn, Pb. We grouped structures sharing the same mapped composition and space group into a single prototype. For example, the compound KGaSb_4 maps to the prototype $a_4b_1d_1_{62}$, where “62” denotes the space group number. This procedure yielded 260 unique prototypes, each representing a distinct combination of composition class and crystallographic symmetry, serving as the basis for our large-scale decoration and creation of hypothetical compounds.

3.1.3 Chemical decoration of prototypes. We applied a systematic chemical decoration strategy to generate hypothetical Zintl phases, using the same element-grouping scheme described above. In this step, we expanded set of elements c to include Si, Ge, Sn, and Pb, as these elements can also exhibit a +2 nominal charge state. During decoration, each element in a prototype was substituted only by elements from the same predefined sets, preserving the nominal oxidation state pattern of the original structure. For example, consider the prototype $a_4b_1d_1_{62}$, which includes four a atoms, one b atom, and one d atom. Given 4 possible elements in a , 5 in b , and 4 in d , this prototype yields $4 \times 5 \times 4 = 80$ unique decorated compositions in the parent prototype structure. Applying these substitutions across all 260 prototypes, we generated a total of 90 856 decorated structures, including the parent ICSD prototypes. After excluding decorations containing Eu and Yb, which are computationally more challenging to model with DFT due to localized f electrons, we retained 72 696 decorated structures for further thermodynamic stability evaluation.

3.2 Upper bound energy minimization (UBEM)

3.2.1 Graph structure and graph neural network. To prepare crystal structures for input to the Graph Neural Network (GNN), we first convert each structure into a graph representation. In this graph, nodes correspond to atomic sites, and edges are defined by connecting each atom to its 12 nearest neighbors based on raw interatomic distances, with periodic boundary conditions taken into account. The node features encode only the elemental identity of each atom (as one hot encoding), while the edge features capture the raw interatomic distances (in Å) between connected atoms. Our GNN architecture employs six message-passing layers, following a design similar to that in ref. 18 and 21. To ensure the model is scale-invariant, we normalize each structure so that the shortest interatomic distance is rescaled to 1 Å, following the approach introduced in ref. 20. This scaling ensures that the model learns a scale-invariant representation of the crystal structures.

3.2.2 Training dataset. We trained the GNN to predict the volume-relaxed total energy of crystal structures, serving as a surrogate model for constrained volume-only DFT relaxations. To construct the training dataset, we selected a representative subset of hypothetical structures described in Section 3.1.3. For computational tractability, we limited our sampling to binary, ternary, and quaternary compositions with fewer than 50 atoms in the unit cell. To ensure chemical and structural coverage, we randomly selected up to 10 compositions per composition type, while ensuring representation across all element types and space groups. This process yielded 6571 structures ($\sim 7\%$ of the full set of 90 856 decorated structures), which were then subjected to DFT volume-only relaxations and used as labeled data for GNN training.

3.2.3 Constrained DFT volume relaxations. Constrained DFT volume relaxations were performed using VASP 5.4.4,^{49,50} interfaced through the Atomic Simulation Environment (ASE) Python package.⁵¹ The volume optimization was carried out in a gradient-free manner using repeated one-shot, self-consistent DFT calculations, similar to our previous work.²¹ We used energy and force convergence thresholds of 1×10^{-6} eV and $0.01 \text{ eV } \text{\AA}^{-1}$, respectively, a plane-wave cutoff energy of 340 eV, and a Γ -centered automatic k -point grid corresponding to 20 divisions along each reciprocal-lattice direction. We employed the Brent method as implemented in SciPy to locate the minimum along the volume–energy curve. To ensure physical relevance and numerical stability, the volume was bounded between $10 \text{ \AA}^3$ (to prevent negative volumes) and twice the volume predicted by the data-mined lattice scheme (DLS), as implemented in Pymatgen.⁵² An initial rough estimate of the volume was first obtained using a linear regression model trained on ICSD structures, prior to applying the DLS predictor.²¹ Structures that hit the upper volume bound during optimization, typically indicating instability due to continuous energy decrease with increasing volume, were pruned from the dataset. Additionally, we filtered out structures with non-smooth volume–energy curves, specifically those where the lowest energy point was more than 10 meV per atom lower than the second-lowest point, to eliminate numerical artifacts. For



Mn-containing compounds, we considered a limited set of randomly chosen 5 different magnetic configurations, including the ferromagnetic (FM) configuration and non-magnetic state. The lowest total energy among these configurations was taken as the constrained volume-relaxed energy. Full computational details are provided in ref. 21 and 53.

3.2.4 Data quality control of training dataset. To ensure the reliability of the training dataset, we implemented several quality control measures to identify and remove problematic or redundant structures from the DFT dataset. Although relaxations were initiated from different prototype structures, we occasionally observed that multiple structures of the composition converged to an identical relaxed configuration. To avoid redundancy, we retained only the lowest-energy structure for each composition among these duplicates. In addition to removing structural duplicates, we excluded any entries with energies or volumes that deviated significantly from expected physical ranges, as these often indicate convergence failures or spurious minima. These filtering steps ensured that the final dataset used for GNN training was both representative and free from artifacts, thereby improving model fidelity.

3.2.5 GNN training. From the 6571 training structures, we employed stratified random sampling to hold out 658 structures for testing. Additionally, we randomly selected 5% of the compositions and held out all associated structures. The model was trained with a batch size of 64 structures over 100 epochs. To optimize the training process, we utilized the AdamW algorithm, starting with an initial learning rate of 10^{-4} , which decayed by 10^{-5} at each update step. The weight decay was initially set to 10^{-5} and similarly decayed by 10^{-5} at each update step. Six dataset sizes were considered to build the learning curve. For each of the six model, we first held out 10% of the structures using stratified sampling for testing, *i.e.*, each composition was sampled proportionately. Subsequently, we calculated the mean absolute error (MAE) of the 10% test sample from the stratified set.

3.2.6 Phase stability analysis. We assess the thermodynamic stability by computing the decomposition energy (E_{decomp}), which is determined from a convex hull analysis by considering competing phases from the ICSD. The total energy of ICSD structures is taken from the NREL Materials Database,²⁴ as described in our prior studies.^{18,21} The total energy calculations in the NREL Materials Database use PBE exchange–correlation functional with PAW pseudopotentials, and energy and force convergence thresholds of 1×10^{-6} eV and 0.01 eV Å⁻¹, respectively. A plane-wave cutoff energy of 340 eV and a Γ -centered automatic k -point grid with 20 divisions along each reciprocal-lattice direction was used. E_{decomp} is a measure of the thermodynamic stability of a phase against decomposition into competing phases. For an unstable phase, E_{decomp} is the minimum energy that the formation energy of that phase has to be lowered (more negative) before it becomes stable and similarly, for a stable phase, it is the minimum energy the formation energy has to be increased (less negative) to render it unstable w.r.t. the convex hull. For each composition, the lowest-energy polymorph is chosen to construct the convex hull and evaluate E_{decomp} . The more negative E_{decomp} is the more stable

(deeper in the convex hull) is the phase. To be self consistent, the convex hulls were reconstructed by considering the volume-relaxed energy predicted with GNN of the initially predicted stable phases from the UBEM approach. After this re-evaluation, the number of predicted stable phases from each approach decreases. The number of newly-predicted phases reported in the Results and discussion is after this re-evaluation procedure.

3.2.7 DFT validation. The newly predicted stable phases are validated with DFT by performing full structural relaxation using the same DFT parameters used for constrained volume relaxations described in Section 3.2.3.

3.3 M3GNet

3.3.1 Structure relaxations using M3GNet. We employed M3GNet (v0.2.4), a graph neural network-based interatomic potential, to perform full structural relaxations of the decorated Zintl structures. The default FIRE optimizer was used, with a total force convergence threshold of 0.1 eV Å⁻¹. The maximum number of ionic steps was set to 4000 for binary, and 8000 for ternary and quaternary structures. M3GNet exhibited robustness: over 99.85% of calculations initiated successfully without errors, and fewer than 0.4% of structures failed to converge within the specified step limits. Calculations were performed using TensorFlow 2.11.0, with cuDNN 8.1.1 and CUDA 11.2, on NVIDIA Tesla V100 GPUs coupled with Intel Xeon Gold 6154 Skylake processors.

3.3.2 Phase stability analysis. We performed convex hull analysis as implemented in Pymatgen^{54,55} to calculate E_{decomp} of decorated structures based on their M3GNet predicted total energies. All competing phases from the Materials Project database were considered since M3GNet was trained on the same database (without including total energy calculated with SCAN functional).²² Conceptually, the convex hull analysis is identical to that defined in Section 3.2.6, but with different total energies. We defined stable structures as those with $E_{\text{d}} \leq 0.0$ eV. We performed the convex hull analysis self-consistently by adding the identified initially predicted stable candidates to re-evaluate E_{d} and assign updated stability, similar to the procedure described in Section 3.2.6.

3.3.3 DFT validation. We performed full structural relaxation with DFT of M3GNet-predicted stable phases. Since M3GNet was trained on the Materials Project dataset (without SCAN functional total energies), we performed DFT calculations using VASP 5.4.4 with the same numerical parameters, pseudopotentials, and Hubbard U values as documented in ref. 56. We considered only ferromagnetic spin configurations for phases with magnetic elements since it is the default spin configurations in Materials Project. Specifically, we used a plane-wave energy cutoff of 520 eV and the energy convergence threshold for electronic self-consistency was set to $5 \times 10^{-5} \times N_{\text{atoms}}$ (in eV), where N_{atoms} is the number of atoms in the simulation cell. The energy convergence threshold for structural relaxation is 10 times the electronic self-consistency energy convergence. The first Brillouin zone is sampled with k -point grids that is determined according to $N_k = 64 \frac{(2\pi)^3}{V}$, where N_k is



the number of k points and V is the volume of the simulation cell. The k -point grid was increased to $100 \frac{(2\pi)^3}{V} 64$ if convergence issues were encountered. An effective Hubbard U value of 3.9 eV was used for Mn.

3.4 Random forest model for predicting formation energy

3.4.1 Dataset and features. A Random Forest (RF) model to predict the formation energy (E_f) of Zintl phases was trained on 3087 stable and unstable structures (1809 from UBE approach, 668 from M3GNet, and 610 phases from the ICSD). The 610 ICSD phases include all polymorphs of a given Zintl composition, which means the dataset includes metastable and unstable structures. The inclusion of both stable, metastable, and unstable structures ensures the model is not biased, as previously demonstrated in ref. 18. Mathematically, E_f is defined as: $E_f = E_{\text{compound}} - \sum_i n_i \mu_i^0$, where E_{compound} is the DFT total energy per formula unit of the compound, and n_i is the number of atoms of element i per formula unit. The reference elemental chemical potentials, μ_i^0 , are the reference values under standard conditions, which are obtained through a global fitting to experimental formation enthalpies, as described in ref. 53. We designed 26 hand-engineered features based on the chemical and physical properties of the constituent elements and the overall structure of the Zintl phases. Details of these features are provided in the SI. The RF regression model training, hyperparameter tuning, and cross validation are performed with the scikit-learn Python library.⁵⁷

3.4.2 Model training and test. The optimized model was obtained through a hyperparameter search aimed at minimizing the 5-fold cross-validated mean absolute error (MAE) in predicting E_f . The following hyperparameters were considered in a grid-type search: (a) the number of decision trees ($n_{\text{estimators}}$), (b) the maximum depth of each decision tree (max_depth), (c) the number of features considered at each split in a decision tree (max_features), (d) the minimum number of samples required to split an internal node (min_samples_split), (e) the minimum number of samples required to be at a leaf node (min_samples_leaf). The tuned hyperparameters are listed in Table S1.

4 Conclusions

Zintl phases represent a rich yet underexplored chemical space for discovering novel functional materials. In this study, we systematically explored this space using a scale-invariant GNN model combined with the UBE approach to accurately predict thermodynamic stability. While stability prediction for hypothetical compounds is inherently challenging, our approach successfully identified 1634 stable Zintl phases as confirmed by DFT calculations. The successful (independent) experimental synthesis of several predicted phases further reinforces the reliability of our approach. We benchmarked UBE against pre-trained M3GNet and found UBE to be more accurate in predicting stable Zintl phases. Together, these two ML frameworks led to the discovery of 2041 new stable Zintl phases, with 260

phases independently identified by both models—highlighting high-confidence candidates. Many of the newly predicted phases feature heavy elements and complex structures, making them especially promising for thermoelectrics.

To bridge the gap between prediction and chemical understanding, we developed an interpretable RF model to relate formation energy to underlying bonding characteristics. Feature engineering guided by chemical intuition, followed by SHAP analysis, revealed that ionic bonding plays a dominant role in stabilizing Zintl phases, regardless of composition or structure. Covalent bonding, while secondary, was shown to enhance or hinder stability depending on its interplay with ionic interactions. By integrating high-throughput prediction with interpretable machine learning, our study not only expands the known chemical space of Zintl phases but also provides a chemically grounded framework for understanding phase stability. This dual capability, discovery with explanation, lays the foundation for more rational and accelerated materials design. Looking forward, our future efforts will focus on property calculations of the predicted phases, particularly thermoelectric properties, and the incorporation of site disorder in predicting stability.

Author contributions

Rinkumoni Chaliha: investigation (equal), data curation (equal), writing (original draft). Manish Kothakonda: investigation (equal), data curation (equal), writing (editing). Cheng-Wei Lee: investigation (equal), data curation (equal), writing (editing). Jeffrey Law: investigation (supporting), writing (editing), supervision. Qian Yang: investigation (supporting), writing (editing). Svilen Bobev: investigation (supporting), writing (editing). Prashun Gorai: conceptualization, investigation (equal), data curation (equal), writing (editing), supervision, project administration.

Conflicts of interest

There are no conflicts to declare.

Data availability

The code and data used to train the GNN and RF models, the main results, and structures presented in this work are available at <https://github.com/prashungorai/papers/tree/main/2025/zintl-search-ubem>.

Supplementary information is available. See DOI: <https://doi.org/10.1039/d5ta06210a>.

Acknowledgements

R. C., M. K. K., C.-W. L., and P. G. acknowledge support from NSF through award DMR-2102409. Q. Y. acknowledges support from NSF through award DMR-2102406. S. B. acknowledges support from the U.S. Department of Energy, Office of Science, Basic Energy Sciences, under Award DE-SC0008885. The research was performed using computational resources



sponsored by the Department of Energy's Office of Energy Efficiency and Renewable Energy and located at the NREL. This work was authored in part by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the US Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government.

References

- 1 D. W. Davies, K. T. Butler, A. J. Jackson, A. Morris, J. M. Frost, J. M. Skelton and A. Walsh, *Chem*, 2016, **1**, 617.
- 2 R. Nesper, *Z. Anorg. Allg. Chem.*, 2014, **640**, 2639–2648.
- 3 S. M. Kauzlarich, *Chem. Mater.*, 2023, **35**, 7355–7362.
- 4 A. Ovchinnikov and S. Bobev, *J. Solid State Chem.*, 2019, **270**, 346–359.
- 5 J. Y. Chan, S. M. Kauzlarich, P. Klavins, R. N. Shelton and D. J. Webb, *Phys. Rev. B*, 1998, **57**, R8103.
- 6 Z. Yuan, D. Dahliah, M. R. Hasan, G. Kassa, A. Pike, S. Quadir, R. Claes, C. Chandler, Y. Xiong, V. Kyveryga, *et al.*, *Joule*, 2024, **8**, 1412–1429.
- 7 K. L. Hodge and J. E. Goldberger, *J. Am. Chem. Soc.*, 2019, **141**, 19969–19972.
- 8 N. Varnava, T. Berry, T. M. McQueen and D. Vanderbilt, *Phys. Rev. B*, 2022, **105**, 235128.
- 9 S. C. Sevov, in *Zintl Phases*, John Wiley & Sons Ltd, 2002, ch. 6, p. 113.
- 10 H. Schäfer, *Annu. Rev. Mater. Sci.*, 1985, **15**, 1.
- 11 P. Gorai, A. Ganose, A. Faghaninia, A. Jain and V. Stevanović, *Mater. Horiz.*, 2020, **7**, 1809–1818.
- 12 Q. Ren, D. Chen, L. Rao, Y. Lun, G. Tang and J. Hong, *J. Mater. Chem. A*, 2024, **12**, 1157–1165.
- 13 C. W. Park and C. Wolverton, *Phys. Rev. Mater.*, 2020, **4**, 063801.
- 14 C. J. Bartel, *J. Mater. Sci.*, 2022, **57**, 10475–10498.
- 15 C. J. Bartel, A. Trewartha, Q. Wang, A. Dunn, A. Jain and G. Ceder, *npj Comput. Mater.*, 2020, **6**, 97.
- 16 J. Riebesell, R. E. A. Goodall, P. Benner, Y. Chiang, B. Deng, A. A. Lee, A. Jain and K. A. Persson, Matbench Discovery—A framework to evaluate machine learning crystal stability predictions, *arXiv*, preprint, 2023, arXiv:2308.14920, DOI: [10.48550/arXiv.2308.14920](https://arxiv.org/abs/2308.14920), <https://arxiv.org/abs/2308.14920>.
- 17 A. Jain, *Curr. Opin. Solid State Mater. Sci.*, 2024, **33**, 101189.
- 18 S. Pandey, J. Qu, V. Stevanović, P. S. John and P. Gorai, *Patterns*, 2021, **2**, 100361.
- 19 J. Gibson, A. Hire and R. G. Hennig, *npj Comput. Mater.*, 2022, **8**, 211.
- 20 K. Pal, C. W. Park, Y. Xia, J. Shen and C. Wolverton, *npj Comput. Mater.*, 2022, **8**, 48.
- 21 J. N. Law, S. Pandey, P. Gorai and P. C. S. John, *JACS Au*, 2023, **3**, 113–123.
- 22 C. Chen and S. P. Ong, *Nat. Comput. Sci.*, 2022, **2**, 718–728.
- 23 A. Belsky, M. Hellenbrandt, V. L. Karen and P. Luksch, *Acta Crystallogr., Sect. B: Struct. Sci.*, 2002, **58**, 364–369.
- 24 N. R. E. Laboratory, NREL Materials Database, <https://materials.nrel.gov>, There is no corresponding record for this reference.
- 25 A. Jain, S. P. Ong, G. Hautier, W. Chen, W. D. Richards, S. Dacek, S. Cholia, D. Gunter, D. Skinner, G. Ceder, *et al.*, *APL Mater.*, 2013, **1**, 011002.
- 26 S. Baranets and S. Bobev, *Mater. Today Adv.*, 2020, **7**, 100094.
- 27 W. Tan, Z. Wu, M. Zhu, J. Shen, T. Zhu, X. Zhao, B. Huang, X.-T. Tao and S.-Q. Xia, *Inorg. Chem.*, 2017, **56**, 10576–10583.
- 28 M. O. Ogunbunmi, S. Baranets, A. B. Childs and S. Bobev, *Dalton Trans.*, 2021, **50**, 9173–9184.
- 29 W. Peng, S. Baranets and S. Bobev, *Crystals*, 2022, **12**, 1467.
- 30 K. Rajput, S. Baranets and S. Bobev, *Chem. Mater.*, 2020, **32**, 9616–9626.
- 31 J. Ishida, S. Iimura and H. Hosono, *Inorg. Chem.*, 2018, **57**, 4997–5003.
- 32 A. Balvanz, S. Baranets and S. Bobev, *Acta Crystallogr., Sect. C: Struct. Chem.*, 2020, **76**, 869–873.
- 33 H. Hirt and H. Deiseroth, *Z. Kristallogr. – New Cryst. Struct.*, 2003, **218**, 5.
- 34 X. Zhang, S. Sun and H. Lei, *Phys. Rev. B*, 2017, **95**, 035209.
- 35 T. Ghosh, M. Dutta, D. Sarkar and K. Biswas, *J. Am. Chem. Soc.*, 2022, **144**, 10099–10118.
- 36 X. Zhang, Q. Liu, Q. Xu, X. Dai and A. Zunger, *J. Am. Chem. Soc.*, 2018, **140**, 13687–13694.
- 37 S. R. Brown, S. M. Kauzlarich, F. Gascoin and G. J. Snyder, *Chem. Mater.*, 2006, **18**, 1873–1877.
- 38 E. S. Toberer, C. A. Cox, S. R. Brown, T. Ikeda, A. F. May, S. M. Kauzlarich and G. J. Snyder, *Adv. Funct. Mater.*, 2008, **18**, 2795–2800.
- 39 E. S. Toberer, S. R. Brown, T. Ikeda, S. M. Kauzlarich and G. J. Snyder, *Appl. Phys. Lett.*, 2008, **93**, 062110.
- 40 J. E. Dobson, *Int. J. Digit. Humanit.*, 2023, **5**, 431–449.
- 41 L. Breiman, *Mach. Learn.*, 2001, **45**, 5–32.
- 42 N. E. Zimmermann and A. Jain, *RSC Adv.*, 2020, **10**, 6063–6081.
- 43 J. Gareth, W. Daniela, H. Trevor and T. Robert, *An Introduction to Statistical Learning: with Applications in R*, Springer, 2013.
- 44 S. M. Lundberg and S.-I. Lee, *Adv. Neural Inf. Process. Syst.*, 2017, **30**, 4768.
- 45 J. Qu, V. Stevanović, E. Ertekin and P. Gorai, *J. Mater. Chem. A*, 2020, **8**, 25306–25315.
- 46 J. E. McGrady, F. Weigend and S. Dehnen, *Chem. Soc. Rev.*, 2022, **51**, 628–649.
- 47 G. A. Papoian and R. Hoffmann, *Angew. Chem., Int. Ed.*, 2000, **39**, 2408–2448.
- 48 S. A. Miller, P. Gorai, B. R. Ortiz, A. Goyal, D. Gao, S. A. Barnett, T. O. Mason, G. J. Snyder, Q. Lv, V. Stevanović and E. S. Toberer, *Chem. Mater.*, 2017, **29**, 2494–2501.
- 49 G. Kresse and J. Furthmüller, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 1996, **54**, 11169.
- 50 G. Kresse and J. Furthmüller, *Comput. Mater. Sci.*, 1996, **6**, 15–50.
- 51 A. H. Larsen, J. J. Mortensen, J. Blomqvist, I. E. Castelli, R. Christensen, M. Dułak, J. Friis, M. N. Groves,



- B. Hammer, C. Hargus, *et al.*, *J. Phys.: Condens. Matter*, 2017, **29**, 273002.
- 52 I.-H. Chu, S. Roychowdhury, D. Han, A. Jain and S. P. Ong, *Comput. Mater. Sci.*, 2018, **146**, 184–192.
- 53 V. Stevanović, S. Lany, X. Zhang and A. Zunger, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 2012, **85**, 115104.
- 54 S. P. Ong, L. Wang, B. Kang and G. Ceder, *Chem. Mater.*, 2008, **20**, 1798–1807.
- 55 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.
- 56 Materials Project, <https://github.com/materialsproject/pymatgen/blob/master/src/pymatgen/io/vasp/MPRelaxSet.yaml>.
- 57 F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot and E. Duchesnay, *J. Mach. Learn. Res.*, 2011, **12**, 2825–2830.

