

ARTICLE OPEN



Interpretable discovery of semiconductors with machine learning

Hitarth Choubisa^{1,5}, Petar Todorović^{1,5}, Joao M. Pina¹, Darshan H. Parmar¹, Ziliang Li¹, Oleksandr Voznyy², Isaac Tamblyn^{3,4} and Edward H. Sargent¹

Machine learning models of material properties accelerate materials discovery, reproducing density functional theory calculated results at a fraction of the cost^{1–6}. To bridge the gap between theory and experiments, machine learning predictions need to be distilled in the form of interpretable chemical rules that can be used by experimentalists. Here we develop a framework to address this gap by combining evolutionary algorithm-powered search with machine-learning surrogate models. We then couple the search results with supervised learning and statistical testing. This strategy enables the efficient search of a materials space while providing interpretable design rules. We demonstrate its effectiveness by developing rules for the design of direct bandgap materials, stable UV emitters, and IR perovskite emitters. Finally, we conclusively show how DARWIN-generated rules are statistically more robust and applicable to a wide range of applications including the design of UV halide perovskites.

npj Computational Materials (2023)9:117; <https://doi.org/10.1038/s41524-023-01066-9>

INTRODUCTION

Inverse materials design, prediction of a structure and composition exhibiting targeted properties, is used to accelerate materials discovery for light emission, sensing, lasing, energy harvesting, and energy storage. Recently, deep learning (DL) models have predicted the properties of molecular and inorganic crystals^{7–9}. However, even with a deep learning acceleration of 10^5 in predicting properties of materials as compared to a single DFT calculation (Supplementary Note 1), exploring the compositional and structural space using existing models remains infeasible: there are $\sim 10^7$ inorganic ternary¹⁰, $\sim 10^{10}$ quaternary compounds^{10,11}, and even more variations for alloyed and multinary compositions¹². Therefore, these property prediction models are usually combined with a search algorithm such as a genetic algorithm (GA) for materials space search^{13–15}. This enables the prediction of materials with the optimal set of properties^{16,17}. However, neither of these components directly enables interpretability and explainability of the behavior of materials and the properties they exhibit.

The development of methods that are adaptable to different applications is equally important. For example, an ML model trained to predict bandgaps of materials can be used to search for materials that emit across a wide range of wavelengths such as UV (<400 nm) or IR (>700 nm). An effective interpretability method should be able to extract design rules for all these applications. Tools such as GNNExplainer¹⁸ explain the origin of candidate material properties, but they are not efficient at extracting chemical rules and theories from the trained property prediction model. Furthermore, such approaches apply to only certain types of ML methods: for instance, GNNExplainer identifies the subgraph of the input graph structure to the GNN that is dominating the prediction by maximizing mutual information between various possible subgraphs and outcome prediction. GNNExplainer is effective at explaining outcomes for a well-trained Graph Neural

Network but cannot be applied to neural networks other than GNNs such as generative models¹⁹.

To overcome the challenge of efficient search and interpretability, we sought to develop a machine-learned framework, one that we term DARWIN: Deep Adaptive Regressive Weighted Intelligent Network. There are three components to DARWIN: a surrogate model, a search algorithm, and the means to distill knowledge in a way that humans can understand. We combine property prediction models and search algorithms with a supervised learning component to extract scientific insights. As part of the approach, we first generate multiple candidates that meet the desired target properties such as stability and UV bandgap. We then use statistical techniques and supervised ML to generate and identify relevant statistically significant chemical rules (Fig. 1 for a summary of the approach). While the approach itself does not make any assumptions about the property prediction model or the search algorithm, we use GNNs as surrogate models and GA as a search algorithm for demonstration.

This paper is organized as follows: we discuss the three components of DARWIN: ML surrogate models that predict material properties using unrelaxed structures, integration of an evolutionary algorithm, and finally, methods for extracting interpretable chemical design rules. We demonstrate the practicality of DARWIN through two use cases: the design of stable UV light-emitting materials and direct bandgap materials.

RESULTS AND DISCUSSION

ML surrogate models

We focus on optoelectronic applications of materials and therefore, train ML models for three properties: energy above the hull, bandgap, and nature of bandgap. Data for energy above the hull and direct–indirect classification was obtained from the Materials Project database^{20,21}. We train the bandgap regressor using a recently published HSE06 xc-functional based dataset²² (refer to

¹Department of Electrical and Computer Engineering, University of Toronto, 10 King's College Road, Toronto, ON M5S 3G4, Canada. ²Department of Physical and Environmental Sciences, University of Toronto, Scarborough, ON M1C 1A4, Canada. ³Department of Physics, University of Ottawa, Ottawa, ON K1N 6N5, Canada. ⁴Vector Institute for Artificial Intelligence, Toronto, ON M5G 1M1, Canada. ⁵These authors contributed equally: Hitarth Choubisa, Petar Todorović. ✉email: isaac.tamblyn@uottawa.ca ted.sargent@utoronto.ca

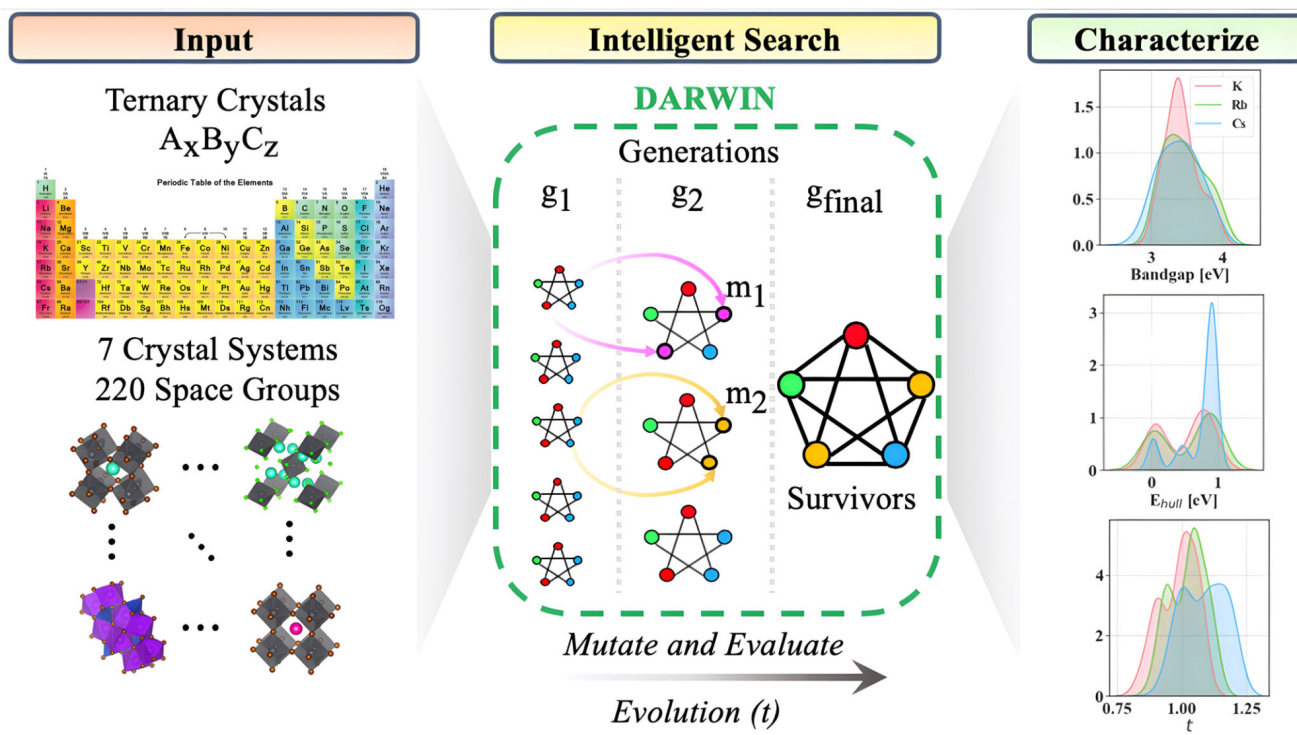


Fig. 1 Different steps of DARWIN. Input crystals were generated using substitutions in prototype structures and spanned over 7 crystal systems and 220 space groups. DARWIN uses trained Graph Networks as surrogate models and mutations to find new candidates that meet the target specifications. Both the negative and positive pools of candidates generated are then characterized using different chemical featurization and subjected to supervised learning combined with statistical testing. The statistically significant rules derived in this way enable the discovery of new compounds and uncover new chemical trends which can be intuitively explained to experimentalists.

“Methods” subsection on “Data generation for ML” for more details; please see Supplementary Fig. 1 for analysis of the data distribution across training, validation, and testing splits of the relevant properties and Supplementary Fig. 2 for distribution of the crystal structures).

We use GCNs for property prediction. GCNs generate a global representation of the crystal structure from the chemical features representative of each element at a given node and edge feature (Fig. 2a, see the “Methods” subsection on “General crystal graph network structure” for details). Several graph convolutional network (GCN) architectures^{6,9} have been reported in the literature to predict properties based on DFT relaxed structures but very few have been reported^{12,23} that predict properties based on unrelaxed structures, which is necessary to perform high-throughput screening without performing computationally expensive DFT geometry optimization.

It is worth noting that, in addition to GCNs, recent progress in predicting material properties has been enabled by the use of generative models^{24,25} such as the invertible crystallographic representation¹⁹ and diffusion-based graph generative model²⁶. Generative models are better at predicting geometry-optimized structures accurately; it remains to be clarified whether they are superior in predicting material properties than the feed-forward ML models. For instance, the most accurate model on Matbench structure-based property prediction challenges is ALIGNN²⁷, which is not a generative model. Herein we explore the use of graph neural networks as property prediction surrogate models. We suggest that these can potentially be replaced with generative models; for these become more accurate without change to the interpretability framework.

To solve this problem, we adapt the MatDeepLearn framework to search, hyperoptimize and benchmark several existing and new GCN architectures (refer to Supplementary Table 1 for all models

considered)²⁸. We also compare the performance of various GCNs against fine-tuning of pre-trained models. We found the most success in learning the map from unrelaxed initial structures to energy above hull and bandgaps through fine-tuning the pre-trained models trained on formation energies obtained from the open quantum materials database (OQMD)²².

Observations of the training experiments for each target property (energy above the hull, bandgaps, and direct/indirect nature of bandgap) are summarized in Supplementary Table 2 and Fig. 2. It can be observed that fine-tuned GCN models outperform other GCNs in predicting energy above the hull, bandgaps, and the nature of bandgaps (direct vs. indirect) from initial structures (Fig. 2). The best GCN model predicts HSE06 bandgaps with mean absolute errors (MAE) of 0.35 eV on test data, and energies above hull with MAE of 0.034 eV/atom (Fig. 2b, c) using unrelaxed structures. Our classifier predicts the direct-indirect nature of the material bandgaps with an F_1 -score of 0.76 and 0.84 using initial and relaxed geometries respectively (Fig. 2d and Supplementary Fig. 3); this value is close to a previously reported (0.89) study on the direct-indirect classification that was limited only to the Kesterite family of compounds²⁹.

Evolutionary algorithm for accelerated search in the chemical space

As the second step of DARWIN, we interface the trained ML models with a search algorithm^{30,31} (evolutionary algorithm/EA) to search through materials space. The fitness function of EA is set as the weighted sum of the mean squared errors of predicted bandgap, energy above the hull, and direct–indirect nature against their desired values. This fitness score is then used to score the candidates. The bottom half is discarded, and the top half is replicated but with each corresponding structure receiving a mutation, generating a new group of candidates to evaluate. We implement the mutation operation as a random elemental

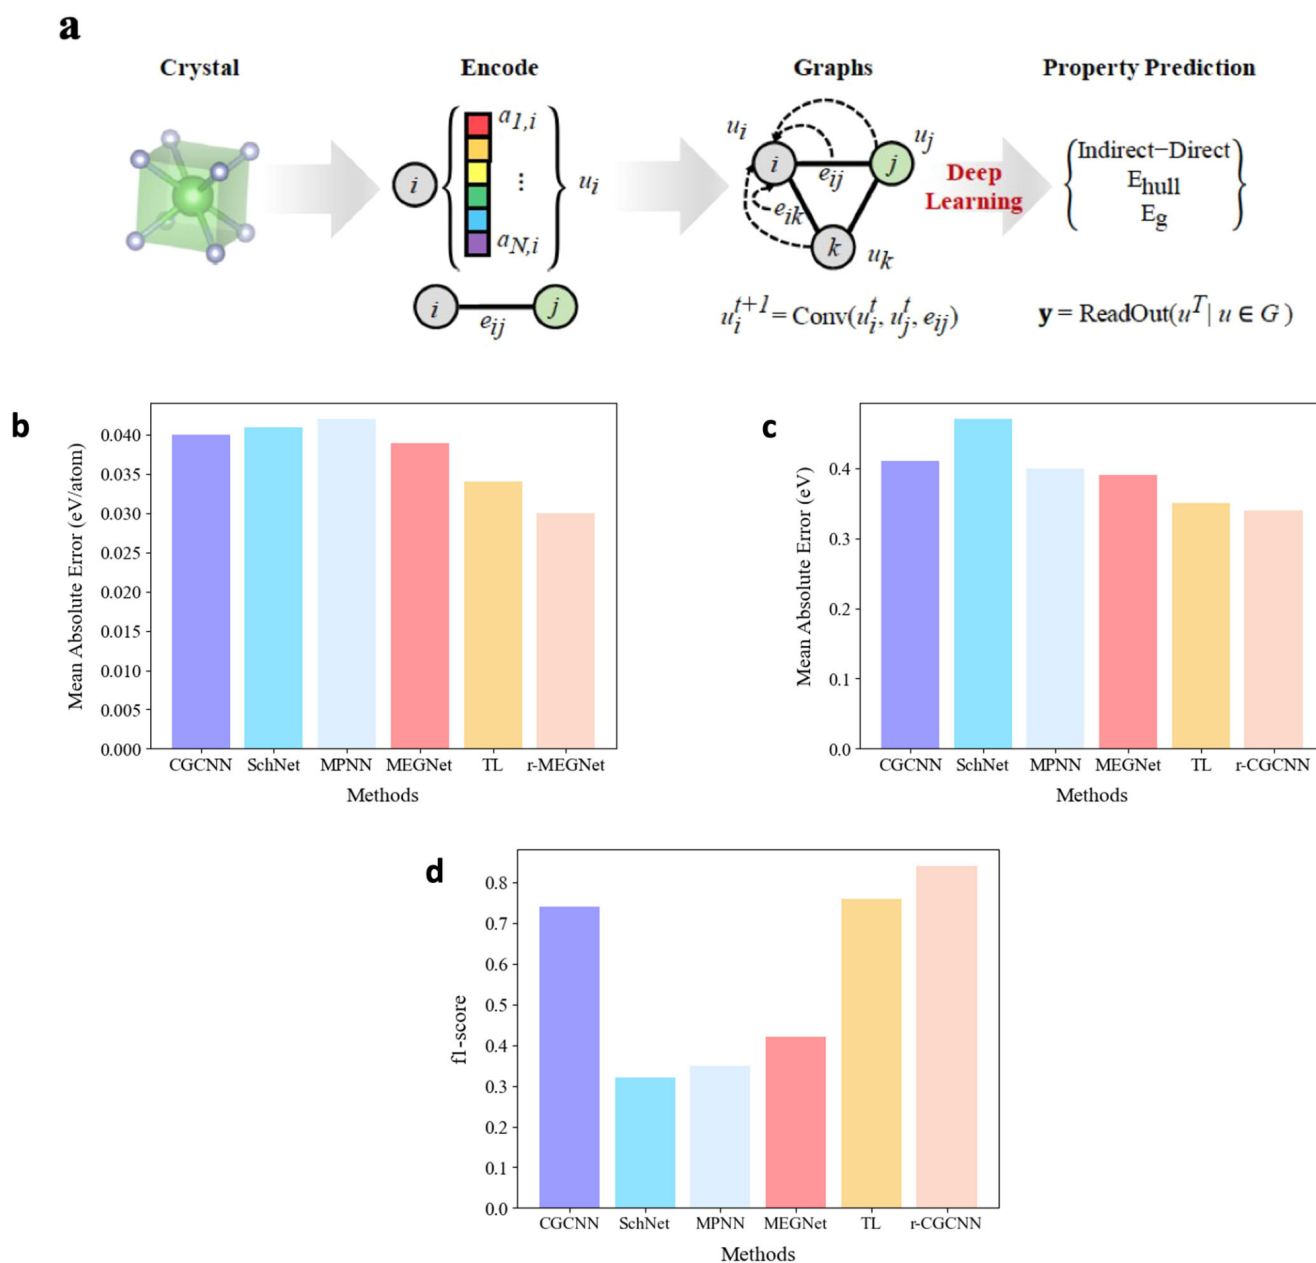


Fig. 2 Surrogate models. **a** Mapping crystals to graph representations through encoding, we train graph neural networks to predict the desired property—bandgaps, energies, and direct/indirect nature. **b** Performance of different GCNs in predicting energies above hull to determine material stability. **c** Performance of different GCNs in predicting HSE06 exchange-correlation functional calculated bandgaps from unrelaxed initial structures without DFT relaxation. **d** Performance of the classifier on direct-indirect classification task using different ML methods (TL: transfer learning approach used in this study; ‘r-MEGNet’ refers to the MEGNet model trained using relaxed optimized geometries as a baseline; ‘r-CGCNN’ refers to the CGCNN model trained using relaxed optimized geometries as a baseline). We only report the best-performing model for relaxed structure predictions out of all those considered (CGCNN, MPNN, SchNet, MEGNet) to setup a robust strong baseline.

substitution of the same oxidation state. This ensures that the charge neutrality of the structure is maintained.

EA relies on our models to predict the properties of interest and evaluate the set of candidates for their fit. Experiments show that mutations alone are enough to direct the search toward the optimal compositions of the large chemical space allowing us to skip crossovers as shown by the decreasing loss as generations of solutions proceeds in time (Fig. 1 for a pictorial representation, “Methods” subsection “Evolutionary algorithm” for more implementation details and Supplementary Fig. 4 for loss as a function of generations).

Interpretability

Although GA combined with a surrogate model can efficiently search the chemical space and lead us to promising candidates, it does not on its own provide an intuitive understanding of the experimental discovery of such materials. The last component of DARWIN solves this problem by identifying chemical features and rules that provide physical insights into the origin of properties that can be consumed by chemists and material scientists in the lab for the design of new materials. All the candidates generated by the GA during its run are collected and categorized into two groups: those that meet the desired target properties and those




Fig. 3 Chemical interpretability. **a** Some relevant and statistically significant (p -value < 0.05) chemical rules generated by DARWIN for the design of stable direct bandgap p -block semiconductors ($E_{hull} < 0.07$ eV/atom). **b** Some of the relevant and statistically significant (p -value < 0.05) chemical rules generated by DARWIN for the design of direct bandgap UV halide-based semiconductors (bandgap range: 3.1 ± 0.3 eV, $E_{hull} < 0.07$ eV/atom). Here, $\mu(\cdot)$ represents mean, $\min(\cdot)$ represents the minimum, $\max(\cdot)$ represents the maximum and $\sigma(\cdot)$ represents the standard deviation of the quantity enclosed. N_v^p is the number of p -valence electrons, T_m is the melting point temperature, C is the column number in the periodic table, Z_m is the Mendeleev number, A is the atomic mass, R_{cov} is the covalent radius, N_v is the number of valence electrons, R is the row number in the periodic table and OEDW is the optimal electronegativity difference window.

that do not. Materials in the two groups are featured using several chemical features and operations on them (Supplementary Note 2 for an exhaustive list of properties and operations) such as the electronegativity difference between B and X site of ternaries ($A_xB_yX_z$), (range, standard deviation, mean, sorted-difference) of electronegativity and elemental chemical properties, HOMO–LUMO corresponding to all the atoms and band centers.

We train a simpler ML model such as Random Forest to learn the classification between the two groups using the generated features. Each feature then acts as a chemical rule and is characterized through two parameters: its relative importance and its statistical significance. We demonstrate two paradigms for acquiring importance: (1) Spearman's coefficient; (2) permutation importance

obtained using Random Forests. Post assignment of importance, we identify the statistical significance of each of these chemical rules using the Kruskal–Wallis H -test (Fig. 1 for the process summary).

Both of these methods are used in tandem to derive scientific insights. In the following two subsections, we use DARWIN to solve two problems in material science: (a) design of direct–indirect bandgap materials and (b) design of stable UV light emitting direct bandgap materials.

Design of direct–indirect bandgap materials

Origin of the direct–indirect nature of bandgap is of fundamental importance for a material's usage in optoelectronics^{32,33} While a

recent study³⁴ tried to explore its origin, it was focused just on binary III–V semiconducting materials.

Here, we extend the criterion and derive chemical rules that explain the origin of the direct–indirect nature of bandgap across all stable *p*-block semiconductors using DARWIN (Fig. 3a). Our approach identifies that semiconductors composed of higher atomic mass *p*-block elements are more likely to exhibit direct bandgap (elements that have smaller melting temperature T_m with larger covalent radius R_{conv} are favorable). Similarly, the more negative the energy of LUMO among individual atomic orbitals constituting the chemical compound and the more the number of *p*-valence electrons on average, the more likely the compound is to exhibit direct bandgap and is stable. The former of these two rules is what has been reported in literature³⁴. We also observe that as the average electronegativities of the elements increase, the material tends to be a stable direct bandgap material. Thus, the chemical insights discovered by DARWIN not only reaffirm one of the previously reported rules but also provide us with yet another statistically significant chemical rule.

Using these design rules, we modify some of the indirect bandgap materials that are widely used in semiconducting and catalytic applications. To test whether DARWIN-derived rules have wider application, we show both cation modification and show mixed anion compounds. We provide a reference that suggests that the synthesis of such compounds may now be feasible³⁵. The results are shown in Table 1.

Design of stable UV light emitting direct bandgap materials

Next, we use DARWIN to solve a slightly more complicated multi-target materials discovery problem: the discovery of stable direct bandgap UV-light emitting materials (3–4 eV), a vast and relatively unexplored³⁶ chemical space^{10,37}. Findings from the interpretability analysis, when the search is limited to perovskites-like structures, reconfirm known predictive descriptors such as the role of A-site and B-site cations in typical perovskite-based crystals for stabilities (Supplementary Figs. 5–7).

When the search is extended to all ternary halide-based compounds, we find several interesting relationships. The features from Fig. 3b are statistically significant (*p*-value < 0.05) and allow us to predict stable UV light-emitting candidates. It is observed that Δ_{X-B} , the difference between the electronegativity (EN) of the B-site (the second most metallic element in the composition) and X-site (most electronegative anion), ranks high and exhibits a small coefficient of variation ($\sigma/\mu < 0.3$). Further analysis revealed that ΔEN_{X-B} is within a narrow range (0.84, 1.5) whenever the material is a stable UV direct bandgap semiconductor. We denote this specific range as the optimal electronegativity difference window (OEDW).

The knowledge of OEDW was then conveyed to the experimental collaborator who combined it with in-lab constraints and factors such as precursor availability, synthesis conditions, and equipment availability. These factors complemented by limited research in K/Cu-based systems at that time made us choose K_2CuX_3 -based systems as ideal and optimal candidates to try experimentally. K_2CuCl_3 and K_2CuBr_3 were experimentally synthesized via spin-coating with an intermediate anti-solvent dripping step^{38,39}. We found that K_2CuCl_3 meets the target specifications with emission below 400 nm (Fig. 5 and Supplementary Fig. 8 for experimental measurements). K_2CuCl_3 has also recently been synthesized independently⁴⁰. Rb_2CuCl_3 satisfies the OEDW criterion and a recent independent report on Rb_2CuCl_3 saw interesting and encouraging results⁴¹. It is worth emphasizing here that instead of relying purely on search results, DARWIN aims to express the predictions in a chemical language that speaks to experimentalists: this enabled us to choose a chemical system that satisfied chemical constraints, leading to UV emission; and enabling experimentalists to incorporate chemical knowledge

such as solubility of precursors, temperature parameters that are otherwise difficult to parameterize, and model using ab-initio methods.

We also performed DFT simulations to verify the predicted optical properties of K_2CuCl_3 . The initial structure was obtained by substituting the prototype structure Eu_2CuS_3 . The initial positions are then relaxed using GGA xc-functional with an energy convergence criterion of 0.0001 eV and a maximum force convergence criterion of 0.01 eV/Å. Simulated and experimental XRD peaks match indicating that the structures obtained after structure optimization is close to the one obtained through experiments (Fig. 4d). Band structures calculations using HSE06 exchange–correlation functional were performed on the relaxed geometry (Fig. 4a). The results from the $E-k$ plot (Fig. 4b) indicate a direct bandgap at the Γ -point. Further analysis of the elemental contributions in the orbitally resolved projected density of states (PDOS) reveals that the halide species significantly contributes to the valence band maxima (VBM) of such materials and the B-cation dominates the conduction band minima (CBM) (Fig. 4b), thus rationalizing the observation that Δ_{X-B} is a good predictor of the bandgap. Specifically, it is observed that in K_2CuCl_3 , K^+ does not contribute to the electronic structure and that the strong orbital interaction of the Cu and Cl species leads to the observed optical properties^{40,42}.

We also used these rules to propose materials that have not been synthesized before i.e., not reported in academic papers nor in materials databases (OQMD, Materials Project, AFLOW). These materials have been compiled in Table 2 with a more comprehensive list added as Supplementary Table 1.

Design of stable IR light emitting direct bandgap perovskites

To test the broader application of the approach, we further apply DARWIN to search for stable direct bandgap IR halide perovskite materials. We focus on a target direct bandgap of 1.2 eV, for this is of interest in tandem solar cells^{43,44}. We initialize the search with halide-based compounds of general formula ABX_3 ($X = Cl, Br$ and I ; refer to Methods subsection on Parameters for genetic algorithm search and optimization of candidates). Results of the interpretability analysis are shown in Fig. 5.

Some of the chemical rules obtained via this analysis simply reconfirm prior literature. For instance, one of the prominent features is $\sum_{\max(R)} T_m$ (T_m : melting temperature; R : row number (1–7) in the periodic table) which shows a negative Spearman correlation and is statistically significant under the Kruskal–Wallis H -test⁴⁵. This indicates that to achieve bandgap with IR emission, the elements must be heavy, and this is supported by existing literature such as iodine-based perovskites including $MAPbI_3$ and $CsPbI_3$ having small bandgaps. The melting point of the metals (T_m) and the number of *p*-valence electrons (N_p^v) appear frequently and are statistically negatively and positively correlated with the ability to emit in the IR, respectively. Thus the transition series metals from the periodic table (rows 4–6 and groups 3–10) are not suitable for IR perovskites. Existing Sn-, Ge-, and Pb-based small bandgap materials agree with this picture^{46,47}. The range of the *p*-valence electrons also is linked to the IR emission behavior in a statistically significant way. It is also worth noting that IR emission is also linked to HOMO–LUMO orbital characteristics of the atomic orbitals constituting the material such that if both originate from *s* or *p* orbitals, it is likely to observe a direct bandgap in the IR regime. Few known compounds with IR bandgaps such as $CsSnI_3$ fit the above criterion. These interpretability guidelines can also be used to modify compounds such that they go closer to emitting in the IR zone. Some of the stoichiometric and alloyed compounds designed following DARWIN interpretation rules are listed in Table 3.

| Table 1. Tuning of indirect bandgap materials to make them direct. | | | |
|--|---------------------------------|---|----------------|
| No. | Indirect materials | Material modulated as per DARWIN-derived rule | Outcome |
| 1 | SiC | Si _{0.5} S _{0.5} C | Direct bandgap |
| 2 | BN | BN _{11/18} Sb _{7/18} | Direct bandgap |
| 3 | Al ₂ Te ₃ | AlInTe ₃ | Direct bandgap |
| 4 | BP | B _{0.67} In _{0.33} P | Direct bandgap |
| 5 | Si ₃ N ₄ | Si ₃ N _{2.5} F _{1.5} | Direct bandgap |
| 6 | InSe | In _{0.75} Tl _{0.25} Se | Direct bandgap |
| 7 | SnS | (SnSb) _{0.5} (ClS) _{0.5} | Direct bandgap |
| 8 | BiS ₂ | BiS _{1.15} Se _{0.5} F _{0.35} | Direct bandgap |
| 9 | BiO ₂ | BiOSe | Direct bandgap |
| 10 | Ge ₃ N ₄ | Ge _{2.06} Sn _{0.94} N ₄ | Direct bandgap |

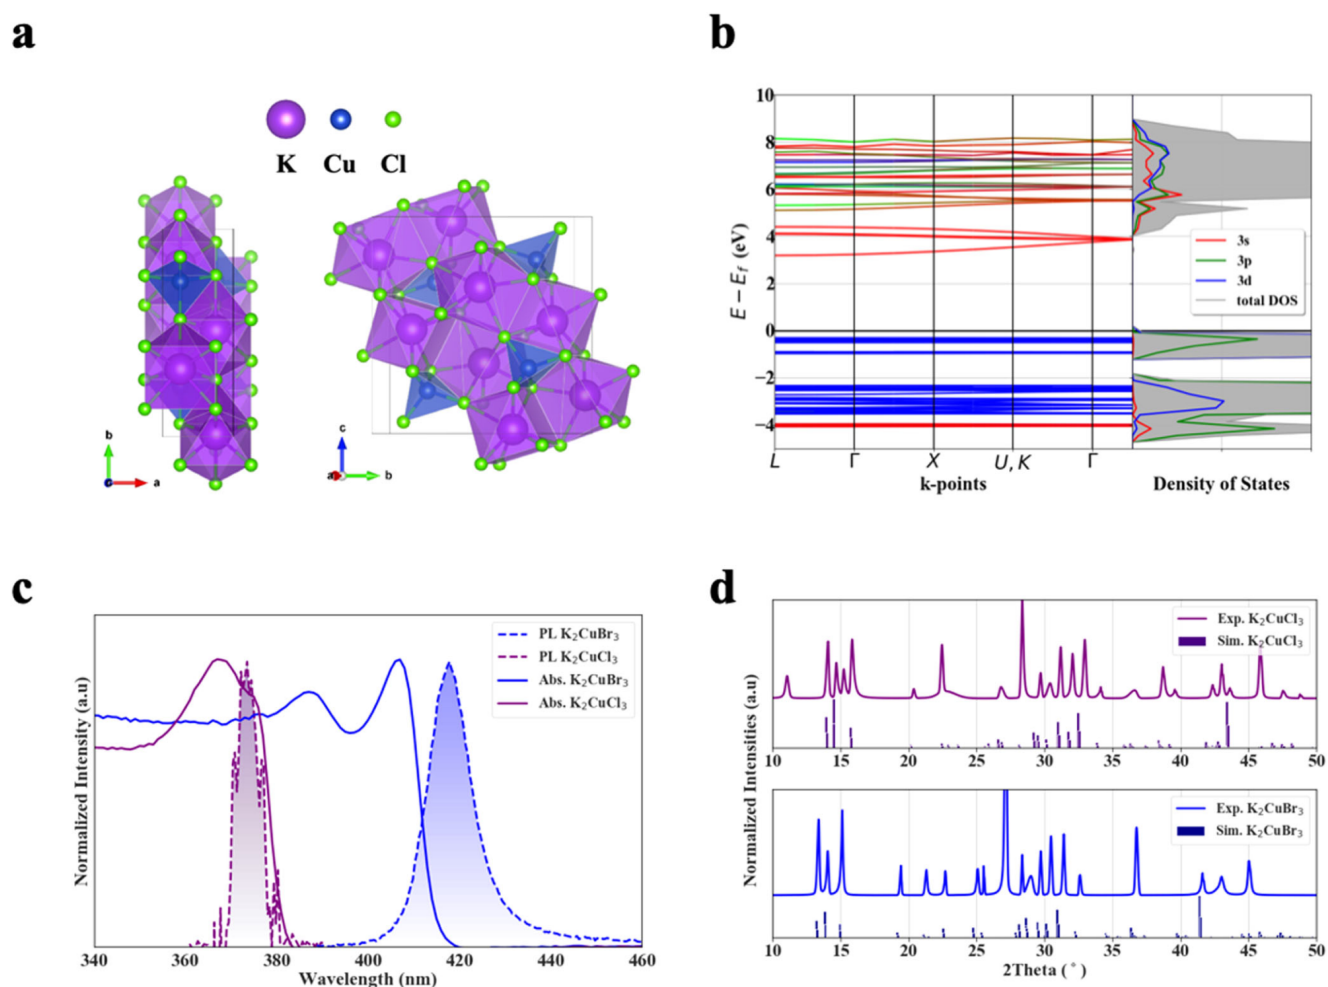


Fig. 4 Experimental realization of K_2CuX_3 and computational studies. **a** Simulated orthorhombic crystal structure ($Pnma$ space group) of K_2CuCl_3 illustrating the 1D chains of $[CuCl_3]^{2-}$ separated by K^+ . **b** Simulated band structure and Density of States of K_2CuCl_3 (Refer to Supplementary Information for similar analysis of K_2CuBr_3). **c** The absorption spectrum and PL profiles of K_2CuBr_3 and K_2CuCl_3 . **d** Simulated²¹ and experimental (powder) X-ray diffraction measurements of K_2CuBr_3 and K_2CuCl_3 .

Ablation experiment

Since there are different components within DARWIN, it is important to ask if all the components are essential for DARWIN's success. We set up an ablation study where we remove the surrogate model and search algorithm with candidates selected from Materials Project. We then compare the performance of this

modified setup with DARWIN in developing interpretable chemical rules for stable UV direct bandgap perovskite materials ($E_g \in [2.8 \text{ eV}, 3.4 \text{ eV}]$) using the correctness and completeness dimensions of interpretability as proposed by Oviedo et al.⁴⁸

Since MP-bandgaps are calculated with PBE xc-functional, they severely underestimate experimental bandgaps. We fit a linear scale between MP-bandgaps and experimental bandgaps⁴⁹ (refer

to Supplementary Note 4 for the equation; Supplementary Fig. 9 for comparison). We then search MP for materials with UV bandgaps using this linear transformation. We end up with 5 compounds that fit the criterion of having energy-above-hull less than 0.07 eV/atom and direct UV bandgap. However, with a 10% margin of error and 95% confidence interval, the recommended sample size for performing significance analysis (calculate p -values) is 97 which is 8 times larger than the number of materials that could be collected from Materials Project (required sample sizes calculated using power test⁵⁰). Thus, removing the search component of DARWIN leads to an inability to calculate statistical significance due to a lack of promising candidates. Even if the significance values obtained were correct, the critical feature we obtained for UV design specifically, OEDW was found to be statistically insignificant (p -value > 0.05) with a very small Spearman correlation of 0.07. On the other hand, OEDW was found to be statistically significant using DARWIN for UV light-emitting perovskite design with a Spearman correlation of 0.54 (Fig. 6a). Both observations indicate that the baseline interpretability approach fails on the completeness front.

Finally, we train a random forest classifier on the collected data from MP. The top non-trivial features obtained are neither statistically nor highly predictive in the larger pool of candidates that were predicted by the search algorithm and surrogate models of DARWIN. Within the top 30 chemical rules that were predicted

by this small dataset, 63% of them turned out to be statistically insignificant (p -value > 0.05) on the larger material candidates generated by the evolutionary algorithm (Fig. 6b). This means that relying just on screening MP, in this case, would limit exploration of candidates and wrong chemical insights. This baseline approach, therefore, fails both on the correctness and completeness front of interpretability. This shows that it is the combination of an accurate surrogate model, large candidate pools generated using a search algorithm, and feature-based interpretability analysis that makes DARWIN effective for interpretable ML for materials discovery. For the application cases shown, the approach not only recovered the known rules for the design of materials but also discovered new chemical rules. These rules enabled the discovery of materials that met the design specifications with human-in-the-loop. The approach, therefore, enables the interpretability of ML-powered material discovery pipelines in addition to just predictions.

METHODS

Data generation for ML

For predicting the stability and optoelectronic properties of the materials, we use DFT calculations to get energy above the hull and bandgaps coupled with the direct/indirect nature of the band structure. We trained GNNs on energy-above-hull and direct–indirect data obtained from the Materials Project on about 117,000 and 45,000 compounds, respectively. The total energy values obtained from the Materials Project are based on the Perdew–Burke–Ernzerhof exchange–correlation functional which has been shown to perform satisfactorily for predicting the stability of the compounds^{6,51}. The list of mp-ids of all the materials data used for this study is attached as part of the SI. We remove all the entries from the dataset that have energy above hull > 2 eV/atom since those represent highly unstable compounds indicating either a very unreasonable geometry or problems with DFT results. To train the bandgap regressor, we use the open-source dataset²² on HSE bandgaps. Furthermore, the direct–indirect classification dataset is unbalanced; therefore, we

| Compound | Ehull (eV/atom) | Bandgap (eV) |
|--------------------------------|-----------------|--------------|
| CsGaCl ₃ | 0.03 | 3.47 |
| CsDyS ₃ | 0.00 | 3.10 |
| SrAuI ₃ | 0.05 | 3.10 |
| Yb ₂ I ₃ | 0.02 | 3.05 |
| TmScN ₃ | 0.04 | 3.11 |




Fig. 5 Interpretation analysis for IR emitting perovskite materials. The analysis shows that p -block elements (in addition to halides themselves) are crucial. This is reflected through both the number of p -valence electrons (N_V^p) as well as the preference for lower melting points (T_m) of the constituent elements. Furthermore, higher electronegativities (χ) also favor direct bandgap IR emission of halide perovskites indicating a shift away from d -block transition elements. Furthermore, orbital character ($c \sim s, p, d, f$) also indicates that homo–lumo that have contributions from s and p -orbitals are more likely to exhibit IR emission.

perform under-sampling and use the balanced subset to train the models. The initial structures, as referred to here, are procured using the MPRester API by specifying the property 'initial_structure'.

General crystal graph network structure

We used the MatDeepLearn²⁸ package combined with PyTorch framework and PyTorch-Geometric module to build and test the crystal graphs and implement the GCN models. The method to encode the crystal structures as graphs has previously been reported in the literature^{8,9,52}. Crystal structures are represented as $G = \{V, E\}$ where V represents atoms represented as nodes and E represents the set of edges connecting two atoms with spatial information. This enables one to represent the 3D geometrical and stoichiometric information of the crystals as graphs. Please refer to Supplementary Note 3 for the exhaustive list of hyperparameters used for the purpose of hyperparameter optimization.

In general, the process can be summarized as follows:

Crystal graphs are fed to the network in batches. We first apply graph convolution operations to them. Convolution operations on a node i can be represented as $\text{Conv}(u_i, u_j^{j \in N(i)}, e_{ij}^{j \in N(i)})$. Several convolution operations have been proposed in literature^{9,52,53}. This quantity is then used to update the node representation for

node i as

$$u_i \rightarrow f\left(u_i, \text{Conv}\left(u_i, u_j^{j \in N(i)}, e_{ij}^{j \in N(i)}\right)\right) \quad (1)$$

This convolution operation is repeated depending on the chosen hyperparameter. Post the convolutions, we perform global pooling of all node features per graph to obtain a fixed-length vector representation of the crystal geometries under inspection (max pooling, min pooling and mean pooling are a few examples). This is represented as

$$U = \text{Pool}_{i \in V}(u_i) \quad (2)$$

In a generalized framework, this vectorial representation is further operated upon using one or more dense layers

$$U \rightarrow \text{act}(A * U + B) \quad (3)$$

where A represents the weights, B represents the biases and act represents the activation function of a dense layer. Finally, the output layer has a single node to enable the prediction of desired properties such as the bandgap or energy above the hull in our case.

In an untrained GCN, the predicted values differ from the ground truth significantly. The error is then backpropagated using a gradient optimizer and all the weights and biases are updated with every epoch till the prediction error reduces to an acceptable level. Changing the hyperparameters such as the number of dense layers, graph convolution layers, pooling type, and optimizer parameters are some of the ways this is usually done.

Table 3. List of promising materials with emissions close to 1.2 eV and not reported in the literature.

| | Final compound | Final bandgap (eV) |
|---|--|---|
| 1 | CsGe _{0.67} Sn _{0.33} I ₃ | 1.56 |
| 2 | CsSn _{0.67} Sb _{0.33} I ₃ | 1.51 |
| 3 | CsGaBr ₃ | 1.16 (CsGaI ₃ has lower stability than CsGaBr ₃) |
| 4 | KGaBr ₃ | 1.17 |
| 5 | CsInBr ₃ | 1.54 |

Transfer learning

Both the architectures (CGCNN and MEGNet) were first fully trained on a dataset of 500k formation energies obtained from OQMD. Post-hyper-parameters optimization, all the convolution and pooling layers were frozen. The number of frozen dense layers, learning rate, and batch sizes were treated as a hyperparameter for transfer learning.

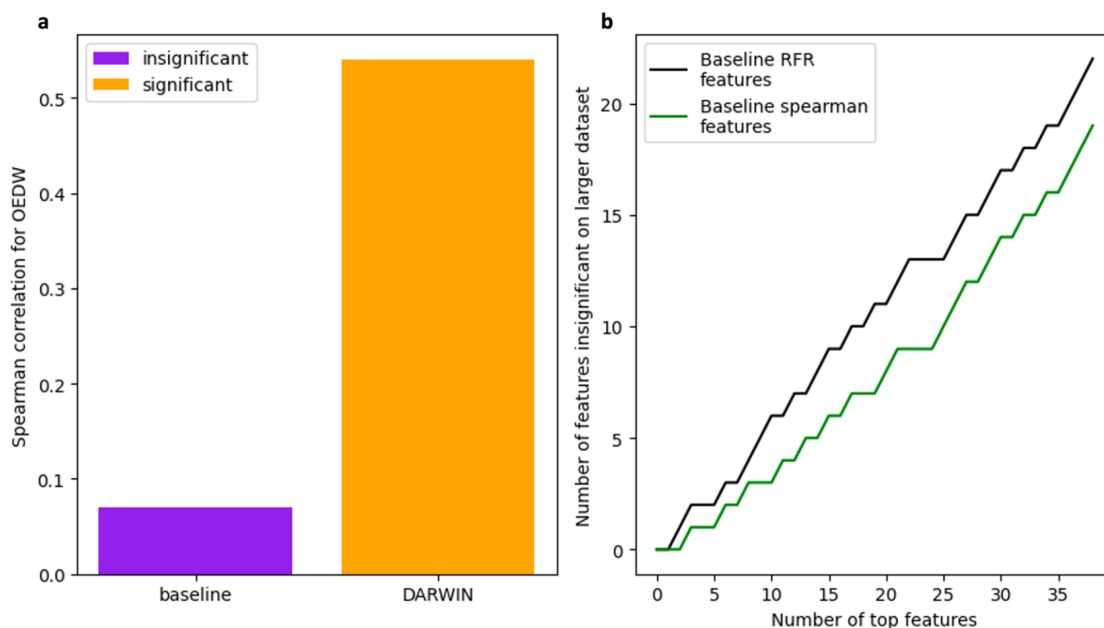


Fig. 6 Comparison of baseline method against DARWIN. **a** When we compare the OEDW interpretability insight on the smaller data that we could extract from MP, we find that not only it turns up a very small Spearman correlation but also statistically insignificant. **b** When we compare the top-ranked chemical rules of the baseline method (either from Spearman analysis or random forest permutation importance), we find that the majority of them turn out to be statistically insignificant (p -value > 0.05) on the larger pool of promising candidates generated by the search algorithm of DARWIN.

Evolutionary algorithm

The EA operates on a surrogate model composed of the three predictive ML models built for the various prediction and classification tasks. A selection criterion is designed for target material properties such as the bandgap value and stability. In general, the multi-step iterative process by which the evolutionary search is implemented is as follows: (1) initialization of primary candidates denoted as the initial generation; (2) prediction of material properties using the ML models; (3) evaluation of the current generation; (4) selection of the fittest candidates; and (5) mutations in the selected individuals, and developing a new generation of candidates.

Over successive iterations, the evolutionary algorithm converges and outputs a set of candidates that are optimal given the current set of selection parameters.

Initialization: In the initialization step, we select a set of elements and generate an initial set of candidates based on the 200 crystal structure types and 7 families. We select the bandgap and energy above the hull which we would like to optimize for and set these search criteria.

Prediction: Crystal graphs are generated via the aforementioned process and fed as inputs into the three pre-trained ML models to obtain prediction values for the bandgap, energy above the hull, and direct–indirect classification.

Evaluation: We evaluate each individual in the current generation given the loss metric as shown in the equation below which is a weighted sum of the squared loss for each individually predicted property and the target selection values, where λ_i are normalizing factors for each loss component. For the selection procedures, we set all the weights to be equal. We initialize with a population of 20 randomly chosen and substituted prototype structures, originally obtained from ALOWlib^{54,55}, set the generation limit threshold at 200.

Selection: Upon evaluating the loss, we rank all individuals by their loss in the current generation and discard the bottom half and retain the remaining population.

Mutation: We then proceed to make a mutation on each top-ranked individual in the population which we define as a single elemental substitution in the crystal structure with the equivalent oxidation state to retain structural charge neutrality. The new set of candidates is then added to the current top-ranked generating a new population and the process is repeated but now starting at evaluation. After multiple iterations, the loss has plateaued, and the EA proposes a set of candidate solutions that ideally match the initial selection criteria. The proposed crystal structures are then aggregated and collected to comprise the candidate solutions for the given target conditions. This process is repeated (100 times in our experiments) for various selection criteria to span the varied bandgap range and design a set of candidate solutions for further analysis and experimental realization.

$$\mathcal{L} = \lambda_1 \left(\hat{E}_{\text{gap}} - E_{\text{gap}}^{\text{target}} \right)^2 + \lambda_2 \left(\hat{E}_{\text{hull}} < E_{\text{hull}}^{\text{target}} \right)^2 + \lambda_3 \left(\hat{E}_{\text{direct}} - 1 \right)^2 \quad (4)$$

Parameters for genetic algorithm search and optimization of candidates

All the case studies shown here were performed with the following set of parameters. The evolutionary algorithm search was conducted for 50 generations with a population size of 50. These evolutionary searches were performed 100 times. The best candidates that meet the threshold requirements were taken as class 1 whereas the worst-performing candidates were labeled as class 0. Performance was measured using the loss function as defined in the “Methods” subsection “Evolutionary algorithm”.

Experimental synthesis—film fabrication

Potassium halide (KX, X = I, Br, Cl), copper halide (CuX, X = I, Br, Cl), dimethylsulfoxide (DMSO), and dimethylformamide (DMF) were purchased from Sigma-Aldrich. Chloroform was purchased from DriSolv. All chemicals were used as received. The precursor solution was prepared by dissolving stoichiometric quantities of KX and CuX in a DMSO/DMF (25/75% v/v) solution (0.5 M) under continuous stirring for 1 h at room temperature. The concentration of the chloride-based precursor solution (in DMSO/DMF 75/25% v/v) was limited to 0.2 M due to the low solubility of the precursors. Glass substrates were O₂ plasma-treated to improve adhesion. The precursor solution was spin-coated onto the substrates via a two-step process: 1000 rpm for 10 s and 3000 rpm. for 60 s. During the second spin step, 0.5 mL of chloroform was poured onto the substrate. The films were then annealed at 110 °C for 10 min. All the samples were prepared in a glove box with an N₂ atmosphere to control the atmospheric conditions.

Material characterization

X-ray diffractograms were recorded using a Rigaku MiniFlex 600 powder X-ray diffractometer equipped with a NaI scintillation counter and using monochromatized Cu K α radiation ($\lambda = 1.5406 \text{ \AA}$). UV–Vis absorption was measured using a Perkin Elmer LAMBDA 950 UV/Vis/NIR spectrometer. PL measurements were collected using a UV–Vis USB 2000+ spectrometer (Ocean Optics). The samples were optically excited using a 355 nm frequency-tripled Nd:YAG laser with a pulse width of 2 ns and a repetition rate of 100 Hz.

There are some missing peaks in the comparison between simulated and powder XRD patterns. We attribute the additional peaks observed in XRD to potentially other remaining lattice planes which exist in the perfect crystal. As with all experimental synthesis, it is possible that certain planes were far favorable in the thin-film fabrication process given the current set of precursor ratios and reaction conditions. This helps explain the mismatch in certain small low-intensity planes. At the same time, it is important to note that neither structures resemble that of the original precursors used to fabricate them. The simulated data assumes that the orientation of the crystals is random; however, powder or thin-film samples would always have preferred orientations. A case in point is the PXRD pattern of K₂CuCl₃ in Figs. 1a and S6 of the study published by Creason et al.⁵⁶. The latter only showed a few peaks compared to the former.

The bandgap and PL of K₂CuX₃ in this paper are different from other published works (Chem. Mater. 32, 6197–6205 (2020); Org. Electron. 86, 105903 (2020)). The crystal structures reported in those works are almost identical to ours—the K₂CuX₃ is composed of 1D [CuX₃]²⁻ chains separated by K⁺. However, their optical characterization is based on single crystals, which may have significant differences in optical properties compared to our solution-processed thin films. Overall, we attribute this difference as a function of the material preparation method which would also cause a discrepancy regarding the bandgap prediction, and so the current deviation is acceptable.

DATA AVAILABILITY

Open-source data was obtained from the Materials Project procured using Pymatgen MPRester API, formation energy data from OQMD, and HSE06 bandgap data from ref. ²².

CODE AVAILABILITY

Model training and hyperparameter optimization was done using MatDeepLearn. All the necessary codes for the analysis can be found at <https://github.com/hitarth64/DARWIN>.

Received: 18 November 2022; Accepted: 9 June 2023;
Published online: 29 June 2023

REFERENCES

- Pilania, G., Wang, C., Jiang, X., Rajasekaran, S. & Ramprasad, R. Accelerating materials property predictions using machine learning. *Sci. Rep.* **3**, 2810 (2013).
- Tabor, D. P. et al. Accelerating the discovery of materials for clean energy in the era of smart automation. *Nat. Rev. Mater.* <https://doi.org/10.1038/s41578-018-0005-z> (2018).
- Schmidt, J., Marques, M. R. G., Botti, S. & Marques, M. A. L. Recent advances and applications of machine learning in solid-state materials science. *npj Comput. Mater.* <https://doi.org/10.1038/s41524-019-0221-0> (2019).
- Kohn, W. & Sham, L. J. Self-consistent equations including exchange and correlation effects. *Phys. Rev.* <https://doi.org/10.1103/PhysRev.140.A1133> (1965).
- Curtarolo, S. et al. The high-throughput highway to computational materials design. *Nat. Mater.* <https://doi.org/10.1038/nmat3568> (2013).
- Park, C. W. & Wolverton, C. Developing an improved crystal graph convolutional neural network framework for accelerated materials discovery. *Phys. Rev. Mater.* <https://doi.org/10.1103/physrevmaterials.4.063801> (2020).
- Gómez-Bombarelli, R. et al. Design of efficient molecular organic light-emitting diodes by a high-throughput virtual screening and experimental approach. *Nat. Mater.* <https://doi.org/10.1038/nmat4717> (2016).
- Duvenaud, D. et al. Convolutional networks on graphs for learning molecular fingerprints. In *Advances in Neural Information Processing Systems*. **28**, (2015).
- Xie, T. & Grossman, J. C. Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties. *Phys. Rev. Lett.* **120**, 145301 (2018).
- Davies, D. W. et al. Computational screening of all stoichiometric inorganic materials. *Chem* <https://doi.org/10.1016/j.chempr.2016.09.010> (2016).
- Isayev, O. et al. Universal fragment descriptors for predicting properties of inorganic crystals. *Nat. Commun.* <https://doi.org/10.1038/ncomms15679> (2017).
- Choubisa, H. et al. Crystal site feature embedding enables exploration of large chemical spaces. *Matter* <https://doi.org/10.1016/j.matt.2020.04.016> (2020).
- Tancret, F. Computational thermodynamics and genetic algorithms to design affordable γ -strengthened nickeliron based superalloys. *Model. Simul. Mater. Sci. Eng.* **20**, 045012 (2012).
- Jensen, J. H. A graph-based genetic algorithm and generative model/Monte Carlo tree search for the exploration of chemical space. *Chem. Sci.* **10**, 3567–3572 (2019).
- Glass, C. W., Oganov, A. R. & Hansen, N. USPEX-evolutionary crystal structure prediction. *Comput. Phys. Commun.* **175**, 713–720 (2006).
- Kim, C., Batra, R., Chen, L., Tran, H. & Ramprasad, R. Polymer design using genetic algorithm and machine learning. *Comput. Mater. Sci.* **186**, 110067 (2021).
- Choudhary, K., Decost, B. & Tavazza, F. Machine learning with force-field-inspired descriptors for materials: fast screening and mapping energy landscape. *Phys. Rev. Mater.* **2**, 083801 (2018).
- Ying, R., Bourgeois, D., You, J., Zitnik, M. & Leskovec, J. GNNExplainer: Generating Explanations for Graph Neural Networks. In *Advances in Neural Information Processing Systems*. **32**, (2019).
- Ren, Z. et al. An invertible crystallographic representation for general inverse design of inorganic crystals with targeted properties. *Matter* **5**, 314–335 (2022).
- Ye, W. et al. Harnessing the Materials Project for machine-learning and accelerated discovery. *MRS Bull.* <https://doi.org/10.1557/mrs.2018.202> (2018).
- Jain, A. et al. Commentary: the materials project: A materials genome approach to accelerating materials innovation. *APL Mater.* **1**, 011002 (2013).
- Kim, S. et al. A band-gap database for semiconducting inorganic materials calculated with hybrid functional. *Sci. Data* **7**, 1–6 (2020).
- Schmidt, J., Pettersson, L., Verdozzi, C., Botti, S. & Marques, M. A. L. Crystal graph attention networks for the prediction of stable materials. *Sci. Adv.* **7**, 7948 (2021).
- Noh, J. et al. Inverse design of solid-state materials via a continuous representation. *Matter* **1**, 1370–1384 (2019).
- Long, T. et al. Constrained crystals deep convolutional generative adversarial network for the inverse design of crystal structures. *npj Comput. Mater.* **7**, 1–7 (2021).
- Xie, T., Fu, X., Ganea, O.-E., Barzilay, R. & Jaakkola, T. Crystal diffusion variational autoencoder for periodic material generation. *Int. Conf. On Learning Representations (ICLR)*, (2022).
- Choudhary, K. & DeCost, B. Atomistic Line Graph Neural Network for improved materials property predictions. *npj Comput. Mater.* **7**, 1–8 (2021).
- Fung, V., Zhang, J., Juarez, E. & Sumpter, B. G. Benchmarking graph neural networks for materials chemistry. *npj Comput. Mater.* **7**, 1–8 (2021).
- Weston, L. & Stampfl, C. Machine learning the band gap properties of kesterite I2-II-V4 quaternary compounds for photovoltaics applications. *Phys. Rev. Mater.* <https://doi.org/10.1103/PhysRevMaterials.2.085407> (2018).
- Padgham, L. & Winikoff, M. *Developing Intelligent Agent Systems: A practical guide* (John Wiley & Sons, 2005).
- Wooldridge, M. *Intelligent Agents: The Key Concepts* https://doi.org/10.1007/3-540-45982-0_1 (2002).
- Soref, R. A. Silicon-based optoelectronics. *Proc. IEEE* **81**, 1687–1706 (1993).
- Yu, P. Y. & Cardona, M. *Fundamentals of Semiconductors Physics and Materials Properties*. (Springer Science & Business Media, 2010).
- Yuan, L. D., Deng, H. X., Li, S. S., Luo, J. W. & Wei, S. H. Unified theory of the direct or indirect bandgap nature of conventional semiconductors. *Phys. Rev. B.* **98**, 245203 (2018).
- Toso, S. et al. Nanocrystals of lead chalcogenides: a series of kinetically trapped metastable nanostructures. *J. Am. Chem. Soc.* **142**, 10198–10211 (2020).
- Tsao, J. Y. et al. Ultrawide-bandgap semiconductors: research opportunities and challenges. *Adv. Electron. Mater.* **4**, 1600501 (2018).
- Allahyari, Z. & Oganov, A. R. Coevolutionary search for optimal materials in the space of all possible compounds. *NPJ Comput. Mater.* <https://doi.org/10.1038/s41524-020-0322-9> (2020).
- Sprigler, B., Schnidrig, S., Todorova, T. & Wild, F. Some thoughts about the single crystal growth of small molecules. *CrystEngComm* <https://doi.org/10.1039/c1ce05624g> (2012).
- Springer Handbook of Crystal Growth* <https://doi.org/10.1007/978-3-540-74761-1> (2010).
- Gao, W. et al. 1D all-inorganic K2CuBr3 with violet emission as efficient X-ray scintillators. *ACS Appl. Electron. Mater.* <https://doi.org/10.1021/acsaem.0c00414> (2020).
- Naewthong, W., Jantapo, W. & Kopwitthaya, A. Synthesis of copper halide nanocrystals and their optical properties. *Nanophotonics and Micro/Nano Optics VII* **11903**, 7–12 (SPIE, 2021).
- Yang, B. et al. Lead-free halide Rb2CuBr3 as sensitive X-ray scintillator. *Adv. Mater.* <https://doi.org/10.1002/adma.201904711> (2019).
- Wang, C., Song, Z., Li, C., Zhao, D. & Yan, Y. Low-bandgap mixed tin-lead Perovskites and their applications in all-Perovskite tandem solar cells. *Adv. Funct. Mater.* **29**, 1808801 (2019).
- Rajagopal, A. et al. Highly efficient perovskite–perovskite tandem solar cells reaching 80% of the theoretical limit in photovoltage. *Adv. Mater.* **29**, 1702140 (2017).
- Kruskal, W. H. & Wallis, W. A. Use of ranks in one-criterion variance analysis. *J. Am. Stat. Assoc.* **47**, 583–621 (1952).
- Ju, M. G., Dai, J., Ma, L. & Zeng, X. C. Lead-free mixed tin and germanium Perovskites for photovoltaic application. *J. Am. Chem. Soc.* **139**, 8038–8043 (2017).
- Wang, W. et al. Highly sensitive low-bandgap perovskite photodetectors with response from ultraviolet to the near-infrared region. *Adv. Funct. Mater.* **27**, 1703953 (2017).
- Oviedo, F., Ferrer, J. L., Buonassisi, T. & Butler, K. T. Interpretable and explainable machine learning for materials science and chemistry. *Acc. Mater. Res.* **3**, 597–607 (2022).
- Zhuo, Y., Mansouri Tehrani, A. & Brgoch, J. Predicting the band gaps of inorganic solids by machine learning. *J. Phys. Chem. Lett.* **9**, 1668–1673 (2018).
- Daniel, W. *Biostatistics: A Foundation for Analysis in the Health Sciences*, 7th edn, 141–142 (Wiley, New York, 1999).
- Schmidt, J. et al. Predicting the thermodynamic stability of solids combining density functional theory and machine learning. *Chem. Mater.* <https://doi.org/10.1021/acs.chemmater.7b00156> (2017).
- Chen, C., Ye, W., Zuo, Y., Zheng, C. & Ong, S. P. Graph networks as a universal machine learning framework for molecules and crystals. *Chem. Mater.* <https://doi.org/10.1021/acs.chemmater.9b01294> (2019).
- Chanussot, L. et al. Open Catalyst 2020 (OC20) dataset and community challenges. *ACS Catal.* **11**, 6059–6072 (2021).
- Mehl, M. J. et al. The AFLOW library of crystallographic prototypes: Part 1. *Comput. Mater. Sci.* <https://doi.org/10.1016/j.commatsci.2017.01.017> (2017).
- Hicks, D. et al. The AFLOW library of crystallographic prototypes: Part 2. *Comput. Mater. Sci.* **161**, S1–S1011 (2019).
- Creason, T. D., McWhorter, T. M., Bell, Z., Du, M. H. & Saporov, B. K2CuX3(X = Cl, Br): all-inorganic lead-free blue emitters with near-unity photoluminescence quantum yield. *Chem. Mater.* **32**, 6197–6205 (2020).

ACKNOWLEDGEMENTS

This work was supported financially by the US Research Center, A Division of Sony Corporation of America (2018 Sony Research Award Program Ref# 2019-0669), and the Natural Sciences and Engineering Research Council (NSERC) of Canada. Authors thank Prof. M. Saidaminov from the University of Victoria for fruitful discussions.

Computations were performed on the SOSCIP Consortium's Niagara and MIST computing platforms. SOSCIP is funded by the Federal Economic Development Agency of Southern Ontario, the Province of Ontario, IBM Canada Ltd., Ontario Centres of Excellence, MITACS, and 15 Ontario academic member institutions. Machine learning models were trained using GPU resources of Northwestern's QUEST computing cluster.

AUTHOR CONTRIBUTIONS

E.H.S. and I.T. supervised the project. H.C., P.T. and I.T. conceived the idea. H.C., P.T., and I.T. performed the ML studies, developed the framework and methodology. J.P. and D.H.P. carried out the experimental fabrication and measurements. H.C., P.T., J.M.P., O.V., I.T., and E.S. discussed the ML results. All authors discussed the results and assisted during manuscript preparation.

COMPETING INTERESTS

The authors declare no competing interests.

ADDITIONAL INFORMATION

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41524-023-01066-9>.

Correspondence and requests for materials should be addressed to Isaac Tamblyn or Edward H. Sargent.

Reprints and permission information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2023

