

# Patterns

## Accurate, interpretable predictions of materials properties within transformer language models

### Highlights

- Text descriptions are efficient in representing materials for property prediction
- Pretrained language models outperform graph neural networks in most cases
- Explanations provided by language models are consistent with human rationales

### Authors

Vadim Korolev, Pavel Protsenko

### Correspondence

korolev@colloid.chem.msu.ru

### In brief

Black-boxed algorithms dominate in materials property prediction. Pretrained large-scale models in conjunction with interpretability techniques counteract the unfavorable tendency by providing clear explanations of suggested outputs.



## Article

# Accurate, interpretable predictions of materials properties within transformer language models

Vadim Korolev<sup>1,2,\*</sup> and Pavel Protsenko<sup>1</sup><sup>1</sup>Department of Chemistry, Lomonosov Moscow State University, 119991 Moscow, Russia<sup>2</sup>Lead contact\*Correspondence: [korolev@colloid.chem.msu.ru](mailto:korolev@colloid.chem.msu.ru)<https://doi.org/10.1016/j.patter.2023.100803>

**THE BIGGER PICTURE** Description, prediction, and explanation are traditionally acknowledged as central aims of science. In the field of materials informatics, the second objective receives the most attention, while the understanding of the resulting structure-property relationships is less emphasized. In this study, we reconcile large-scale language models and human-readable descriptions of crystal structure to facilitate materials design insights. The presented approach surpasses the state of the art in property prediction and provides transparency in the machinery of artificial-intelligence algorithms, thereby possibly improving the trust of materials scientists. In addition, the clarity of text-based representation and maturity of associated explainability methods make the approach appealing for educational uses.



**Proof-of-Concept:** Data science output has been formulated, implemented, and tested for one domain/problem

## SUMMARY

Property prediction accuracy has long been a key parameter of machine learning in materials informatics. Accordingly, advanced models showing state-of-the-art performance turn into highly parameterized black boxes missing interpretability. Here, we present an elegant way to make their reasoning transparent. Human-readable text-based descriptions automatically generated within a suite of open-source tools are proposed as materials representation. Transformer language models pretrained on 2 million peer-reviewed articles take as input well-known terms such as chemical composition, crystal symmetry, and site geometry. Our approach outperforms crystal graph networks by classifying four out of five analyzed properties if one considers all available reference data. Moreover, fine-tuned text-based models show high accuracy in the ultra-small data limit. Explanations of their internal machinery are produced using local interpretability techniques and are faithful and consistent with domain expert rationales. This language-centric framework makes accurate property predictions accessible to people without artificial-intelligence expertise.

## INTRODUCTION

Artificial intelligence (AI) is increasingly perceived as the fourth pillar of modern science<sup>1</sup> rather than a tool complementary to the previous three, namely experiment, theory, and simulation. In the materials science realm, a data-driven approach has been successfully employed to capture complex structure-property relationships. In particular, AI techniques have pushed the frontiers of high-throughput computational screening,<sup>2–5</sup> inverse materials design,<sup>6–9</sup> interatomic potential development,<sup>10–12</sup> and crystal structure prediction.<sup>13–15</sup> Users demand that supervised machine-learning (ML) models involved in the above tasks be accurate first. Driven by this demand, AI practitioners among mate-

rials scientists have developed increasingly complex models by elaborating data representations. A retrospective view of the evolution of graph neural networks in the field<sup>16</sup> should serve as a definitive example: the models that hold global state attributes<sup>17</sup> and many-body interactions<sup>18</sup> have shown growing improvement in performance relative to neural networks trained on a compact set of node and edge features.<sup>19</sup> The state-of-the-art architectures aimed at materials property predictions may contain many thousands or millions of trainable weights (in other tasks, billions<sup>20</sup> and trillions<sup>21</sup>), making a good understanding of internal machinery intractable for human beings. Unsurprisingly, the explainability problem gives rise to distrust and hinders wider applications of ML algorithms. It should be noted that there are



special scenarios where the use of uninterpretable black boxes may be appropriate.<sup>22</sup> Nevertheless, the explanation of model reasoning (“Why is it the answer?”) is a desideratum of effective AI in general<sup>23</sup> and specific<sup>24</sup> contexts.

Explainable AI<sup>25,26</sup> (XAI) is an umbrella term for algorithms intended to make their decisions transparent by providing human-understandable explanations. Following the proposed taxonomies,<sup>27–29</sup> one can differentiate XAI methods based on multifaceted but nonorthogonal dualities: model-specific vs. model-agnostic, intrinsic vs. *post hoc*, and local vs. global explainers. Despite the impressive diversity, only a few XAI approaches are applied in materials science actively. The most remarkable techniques are examined below; for more details, please refer to recent reviews.<sup>30–33</sup> First, we would like to highlight supervised ML algorithms that have inherent transparency. Linear regression models and their extension, generalized additive models,<sup>34,35</sup> provide weight coefficients as an importance metric of relevant features.<sup>36–39</sup> Decision trees<sup>40</sup> are another method that has shown off-the-shelf transparency. For example, probability values in terminal nodes of classification models reveal the combinations of splitting criteria leading to preferable output.<sup>41–43</sup> The next approach involves the derivation of analytical expressions of structure-property relationships. Symbol regression<sup>44</sup> and compressed sensing methods such as least absolute shrinkage and selection operator<sup>45</sup> (LASSO) and sure independence screening and sparsifying operator<sup>46</sup> (SISSO) make it possible to access sets of solutions competitive in terms of accuracy and complexity.<sup>47–53</sup> Finally, there are *post hoc* local explainers that quantify feature importance levels when analyzing opaque ML models. Among such XAI schemes, the Shapley Additive Explanations<sup>54</sup> (SHAP) suite seems to be dominant in materials science applications.<sup>55–66</sup>

The common thread in most of the explainability-aware studies mentioned above relates to employing low-dimensional hand-crafted features as input to a training model. Besides the notorious tradeoff between accuracy and explainability of ML algorithms,<sup>67</sup> a similar compromise is seen in materials representations. Going beyond simplistic physicochemical descriptors, more advanced featurization schemes (for instance, physics-inspired schemes<sup>68</sup>) can hardly be interpreted in terms familiar to domain specialists. Nonetheless, there is an alternative to tabular-like representations; as a consequence, other XAI techniques may come into play. Graph neural networks rooted in deep geometric learning<sup>69,70</sup> successfully cope with data irregularities. In particular, periodic atomistic systems can be processed in a natural way if the concept of a crystal graph<sup>17–19,71–75</sup> is introduced. The propagation and aggregation of information contained in node and edge attributes via message passing and pooling operators are aimed at describing pairwise and higher-order interactions. To date, graph neural networks have archived state-of-the-art performance in predicting a plethora of structure-property relationships.<sup>16,76</sup> The uncovering of such black-box models in order to gain chemical insights is addressed by the development of graph-specific XAI techniques<sup>77–79</sup> and by the adaptation of methods from other domains.<sup>80–82</sup> The field is in its infancy, and therefore explainability-aware studies exploring graph neural networks in materials science are sparse and few.<sup>83–85</sup>

There is no community consensus on defining explainability owing to the diversity in XAI approaches and problems being

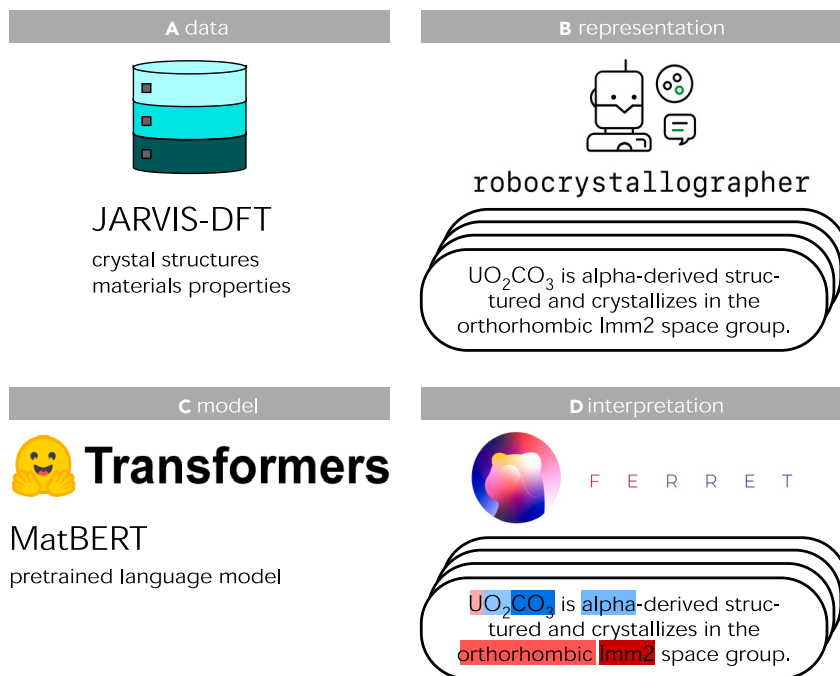
solved. Nevertheless, attempts to specify the term<sup>23,86–88</sup> are united by the idea that the perceiver (domain expert) is as important as the explainer (XAI algorithm). Moreover, cognitive abilities of the former limit model understanding,<sup>89</sup> which is the primary reason why researchers are forced to consider simple features in supervised ML if model reasoning is simple in the first place. We stress that natural language representation of materials is an optimal way to archive interpretability by human beings. The corresponding AI field, natural language processing,<sup>90</sup> has already found successful applications in materials science such as named entity recognition<sup>91–93</sup> and paragraph classification.<sup>92,94,95</sup> At the same time, the potential of natural language features in materials property prediction is fully unexplored to the best of our knowledge.

In this study, we present a language-centric framework able to reconcile high accuracy and interpretability of the prediction of materials properties. Attention-based neural networks trained on text descriptions are thoroughly compared with graph neural networks, including compositional and structure-aware architectures. A classic ML algorithm, random forest, built on force-field-inspired descriptors is included in the benchmark as well. We demonstrate remarkable scalability of the language models that allow state-of-the-art performance to be achieved in a small-data regime. In certain cases, transformers trained on human-readable features surpass graph neural networks despite training-dataset size. The interpretability of our approach is estimated in terms of faithfulness and plausibility. As the analysis showed, XAI approaches can generate sufficient and comprehensible explanations that are consistent with expert decision making.

## RESULTS

### General workflow

The general workflow of the study is presented in Figure 1. The first three stages are the preparation of the main supervised ML components<sup>96</sup>: an input dataset, feature representation, and a predictive algorithm. We start by considering the crystal structures and corresponding property values taken from the Joint Automated Repository for Various Integrated Simulations<sup>97</sup> (JARVIS). The following diverse set of endpoints is taken into account: energy above the convex hull, the magnetic moment, a band gap, spectroscopic limited maximum efficiency (SLME), and topological spin-orbit spillage. All the data were originally obtained at the density functional theory (DFT) level in accordance with standardized calculation procedures. High-throughput computational databases<sup>98–100</sup> such as JARVIS-DFT help investigators to focus on ML tasks (e.g., feature and algorithm benchmarking and model fine-tuning) and to simplify data preprocessing. At the second step, we for the first time implement human-readable text descriptions as materials representation for supervised ML. We use the Robocrystallographer library<sup>101</sup> to automatically generate the proposed representation for thousands of crystal structures, but a similar content written by a human crystallographer is also acceptable. At the third step, advanced language models, namely transformers,<sup>93,102,103</sup> are utilized to extract structure-property relationships based on text descriptions generated at the previous step. All the prediction tasks are examined in the classification mode. Specifically, initially continuous quantities are converted into binary labels based on specific threshold values (see the



**Figure 1. An overview of the language-centric approach**

(A) Initial data are taken from an open computational database containing crystal structures and a diverse set of physical properties calculated at the level of density functional theory (DFT).  
(B) Proceeding from crystal structures, we generate text-based descriptions via an automatic toolkit. Local, semilocal, and global environment features are taken into account.  
(C) Neural networks capable of handling a natural language are trained on the text-based descriptions to classify materials.  
(D) *Post hoc* explainability techniques help to rationalize algorithm decisions at the level of tokens.

section [experimental procedures](#)) to simplify further XAI analysis and comparison with human rationales, which are representable in dualistic terms as well. The last step of our workflow is designed to assess interpretability of the language-centric approach. The trained models are processed within the suite of *post hoc* local XAI techniques.<sup>104</sup>

### Model performance

Aside from the input representation and network architecture, several other factors directly influence language model performance. First, various vocabularies can be applied during tokenization, which is a procedure of splitting text into elementary bits of information. Second, the model can be optionally pretrained on a large text corpus before use in downstream tasks. We present results of training of three models differing in the above characteristics. The transformer models combined with a general and domain-specific tokenizer without pretraining are designated as Bidirectional Encoder Representations from Transformers (BERT) and BERT-domain, after the name of the base architecture.<sup>103</sup> The transformer model pretrained on the corpus of materials science papers<sup>93</sup> (MatBERT) and combined with the domain-specific tokenizer is tested as well. We would like to integrate the presented models into the landscape of modern ML models predicting materials properties. To do so, the following algorithms are included in the benchmark for a comparison with the models mentioned above. Random forest<sup>105</sup> trained on classical force-field-inspired descriptors<sup>106</sup> (RF-CFID) serves as a representative classic ML algorithm built on tabular features. A deep neural network designated as representation learning from stoichiometry<sup>107</sup> (Roost) typifies advanced models trained on chemical compositions. Finally, Atomistic Line Graph Neural Network<sup>18</sup> (ALIGNN) is examined as a state-of-the-art predictor of materials properties.

As a target metric, we calculate the Matthews correlation coefficient (MCC). The metric is recognized as a reliable statistical measure and is preferable to other binary classification metrics, including accuracy and the F1 score.<sup>108</sup> MCCs for all the above models and endpoints are provided in [Table 1](#). MatBERT works surprisingly well, manifesting state-of-the-art performance

in four cases out of five. ALIGNN has the highest MCC only on magnetic or nonmagnetic classification. The overall MCC across all endpoints equals 0.74 and 0.72 for MatBERT and ALIGNN, respectively. RF-CFID and Roost show worse performance than ALIGNN, with one exception. Roost has the second-highest MCC on energy above the convex hull. The same trends are observed for accuracy ([Table S1](#)) and the F1 score ([Table S2](#)). The relatively high efficiency of the structure-agnostic model (Roost) contradicts previous results. Bartel et al. have demonstrated significant improvement in stability predictions owing to inclusion of crystal structure in representation.<sup>109</sup> On the other hand, it is important to note that another endpoint (decomposition energy) has been addressed by those researchers. The interplay of model architecture and the thermodynamic stability criterion seems to be a promising avenue of future work and is beyond the scope of this study.

We can speculate that the superiority of the presented language-centric approach over others rests on knowledge enrichment rather than the model architecture. For instance, MatBERT significantly outperforms BERT-domain (an overall MCC of 0.67), whereas the only difference between them is pretraining of the former. In particular, BERT-domain was trained on 1.6 million to 10.1 million tokens depending on an endpoint. MatBERT was trained on 8.8 billion<sup>93</sup> plus 1.6 million to 10.1 million tokens, taking into account the masked language modeling of the original model. Therefore, the effective dataset size increases by approximately three to four orders of magnitude via fine-tuning.

The approach applied herein—fine-tuning—belongs to the transfer learning paradigm. In materials science, models are usually pretrained on labeled data in a supervised manner.<sup>110,111</sup> By contrast, language models, such as transformers, provide a great opportunity to capture domain knowledge within self-supervised learning. The above-mentioned masked language modeling is a vivid representative of such techniques. The

**Table 1. Model performance in terms of the MCC**

	Energy above hull	Magnetic moment	Band gap	SLME	Spin-orbit spillage
RF-CFID	0.791 ± 0.012	0.735 ± 0.012	0.800 ± 0.013	0.595 ± 0.018	0.492 ± 0.027
Roost	0.885 ± 0.005 <sup>#</sup>	0.762 ± 0.009	0.794 ± 0.020	0.580 ± 0.019	0.482 ± 0.025
ALIGNN	0.878 ± 0.010	0.793 ± 0.009 <sup>*</sup>	0.827 ± 0.011 <sup>#</sup>	0.615 ± 0.027 <sup>#</sup>	0.507 ± 0.026 <sup>#</sup>
BERT	0.788 ± 0.011	0.674 ± 0.014	0.747 ± 0.014	0.446 ± 0.026	0.401 ± 0.027
BERT-domain	0.841 ± 0.013	0.727 ± 0.011	0.791 ± 0.011	0.52 ± 0.04	0.464 ± 0.026
MatBERT	0.901 ± 0.005 <sup>*</sup>	0.788 ± 0.007 <sup>#</sup>	0.845 ± 0.011 <sup>*</sup>	0.629 ± 0.017 <sup>*</sup>	0.519 ± 0.022 <sup>*</sup>

The best coefficient for each endpoint is indicated by an asterisk; the second-best one is indicated by a superscript hash symbol.

proposed language-centric approach allows us to profitably incorporate a massive source of scientific knowledge (journal publications) into a workflow intended for materials property predictions. Henceforth, domain-specific text corpora should be seen as an alternative to high-throughput DFT databases<sup>97–100</sup> containing thousands of crystal structures and calculated properties in the context of transfer learning. Moreover, predictors based on a text description can benefit from processing of large samples of papers even without the training of language models on them. BERT-domain clearly surpasses BERT in terms of the MCC, accuracy, and F1 score (Tables 1, S1, and S2) for all endpoints, exclusively due to the domain-specific tokenizer trained on 2 million materials science articles. This result proves the importance of using a domain-specific vocabulary because the general-purpose BERT tokenizer is unable to meaningfully process chemical formulas, space group symbols, and other data. Another remarkable feature of such preprocessing is that vocabulary construction, even on large corpora, takes negligible computing time. By analogy with corpora of academic texts and computational databases, we view domain-specific tokenization as a low-cost alternative to fine-tuning within a general-purpose vocabulary.

It is well known that ML model performance depends on training-dataset size.<sup>112</sup> Moreover, predictive algorithms significantly differ in sensitivity to the growth of available data.<sup>113,114</sup> To clarify this issue for the considered endpoints, we retrain two best-performing models, MatBERT and ALIGNN, on a part of the original training datasets. Figure 2 shows a strong linear dependency ( $R^2 > 0.99$ ) for a logarithmic scale. MatBERT is more accurate in the ultra-small data limit (it has a systematically higher intercept value in the linear equations), in good agreement with the hypothesis that fine-tuning affords an increase in effective dataset size. Assuming that the linear approximation holds true as the training-dataset size grows, we identify limits of superiority of MatBERT over ALIGNN (Figure 2F). ALIGNN has the highest MCC with a training-dataset size exceeding 16,297, 12,974, and 55,533 entities in the case of the magnetic moment, SLME, and spillage, respectively. On the contrary, MatBERT dominates in classification of energy above the convex hull and a band gap regardless of data availability (it has simultaneously higher intercept and slope values in the linear equations). This is a surprising result because, generally, higher scalability of a graph neural network is expected for the following reason. The ALIGNN architecture incorporates information-rich structure representation, which is capable of absorbing subtle crystallographic features in contrast to human-readable features implemented in MatBERT. Consequently, each additional training point should

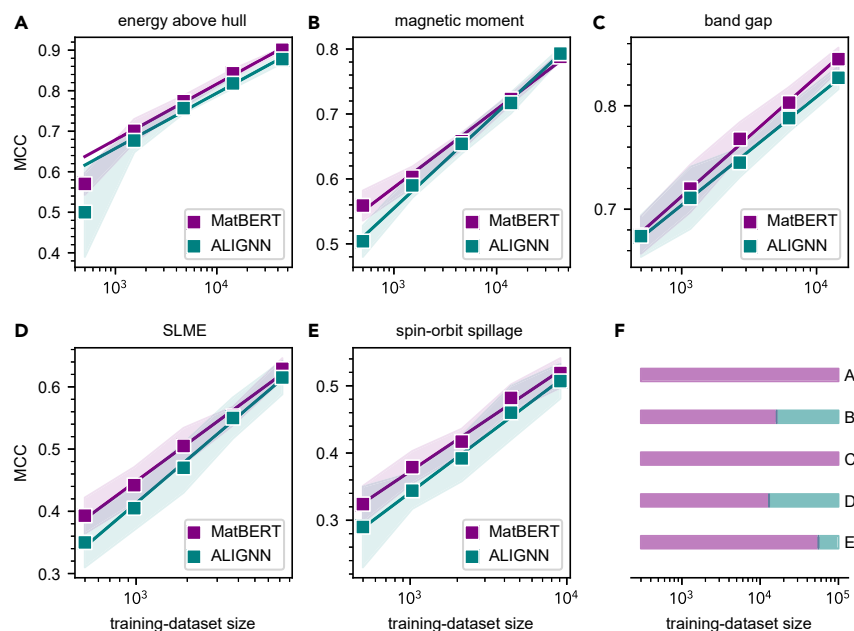
potentially enrich the graph neural network more than the language model. Regardless of whether the state-of-the-art performance of the presented approach is confirmed on larger training datasets, MatBERT shows unexpectedly good performance at the scale of thousands of training samples. As outlined above, input features for language models were generated with Robocrystallographer in a high-throughput manner. Further attempts may be made to clarify how trained models can withstand unperceivable changes (in text descriptions), which may occur when input data are created by human experts. Mature techniques for robustness evaluation of general-domain language models, such as adversarial attacks,<sup>115</sup> are applicable here.

### Model explanation

Now we are going to evaluate interpretability of the language-centric approach. According to the prominent framework introduced in DeYoung et al.,<sup>116</sup> two peculiar aspects are taken into account. First, we examine the ability of XAI techniques to correctly reflect internal machinery of a predictor (i.e., faithfulness). Second, the consistence of explainer output within human reasoning, also called plausibility, is evaluated. We would like to emphasize that the proposed materials representation (entity-level input into a predictor) is a human-readable text description separated into a sequence of tokens. To explain decisions of the black-box model at the level of distinct entities, local *post hoc* XAI techniques are employed. The choice of feature importance methods is explained by the resemblance between how such approaches represent ML reasoning and how human beings tend to perceive a natural language, by highlighting the most meaningful parts of a text. Four techniques are implemented to identify tokens most impactful on the prediction: saliency map extraction via computation of input gradients<sup>117</sup> (hereinafter referred to as saliency maps [SMs]), integrated gradients<sup>118</sup> (IGs), local-interpretable-model-agnostic explanations<sup>119</sup> (LIMEs), and SHAP.<sup>54</sup> We apply several explainers to achieve suboptimal results for a specific task because there is no *a priori* knowledge about which explainer shows state-of-the-art performance.

The faithfulness of XAI techniques is determined via an erasure procedure,<sup>120</sup> which comprises removing some tokens and identifying changes in model confidence. We estimate two faithfulness measures. Specifically, comprehensiveness as an evaluation metric stands for model degradation caused by eliminating the most influential tokens; a larger value is better. On the other hand, sufficiency hinges on model stability if only influential tokens are taken into account; a smaller value is better. The tokens are excluded within a ranking produced by the explainer being analyzed. A more detailed description of both metrics is provided





**Figure 2. Scalability of ALIGNN and MatBERT**

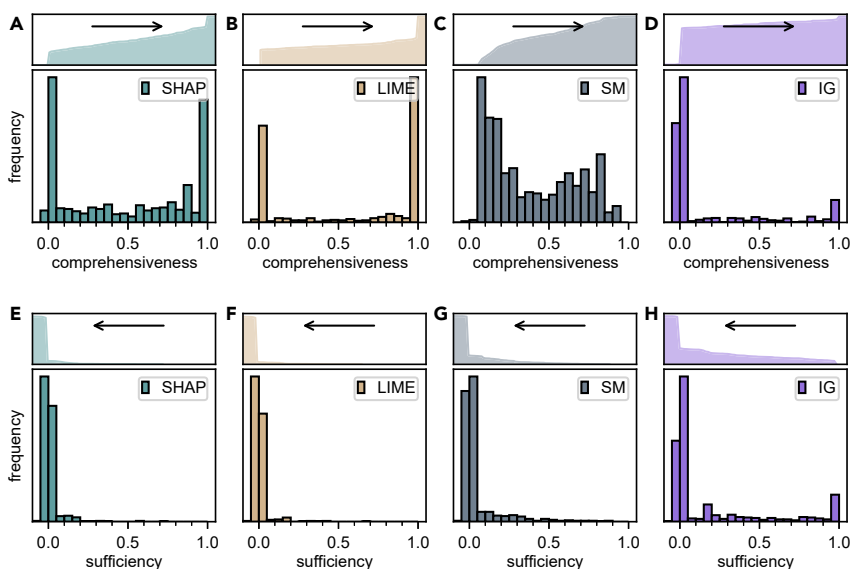
(A–E) Matthews correlation coefficients (MCCs) are shown as a function of training-dataset size for the following classification tasks: (A) energy above the convex hull, (B) the magnetic moment, (C) a band gap, (D) spectroscopic limited maximum efficiency (SLME), and (E) topological spin-orbit spillage. The solid lines denote linear fitting of the data; the smallest dataset (499 entities) in the energy-above-the-convex-hull task is not included in the analysis because of a severe deviation from linear behavior. The shaded areas are standard deviation across 10-fold cross-validation.

(F) Regions of dominance of two examined models in terms of the MCC are marked by the corresponding color.

The tokens highly ranked by the explainer with confirmed faithfulness may be helpful for explaining how a language model processes human-readable materials descriptions. We select the top 5% of nonunique tokens within the ranking given by LIME.

in the section [experimental procedures](#). In the following explainability analysis, we limit our attention to the MatBERT model used in band gap classification (Figure 3). With comprehensiveness, most explainers result in a bimodal distribution with peaks at zero and one. Therefore, there are two distinct groups of test examples differing in the ability of the XAI methods to extract meaningful rationales behind model reasoning. Entities with the score close to one are almost comprehensively interpreted within the corresponding explainer; the opposite is true for the second group of materials descriptions. LIME shows the largest proportion of test examples with high comprehensiveness (>0.9): 49%. SHAP has the second-highest ratio of 25%, whereas SM and IG yield only 2% and 8%, respectively. On the other hand, 8% of test samples exhibit low comprehensiveness (<0.1) across all implemented explainers. This set is distinguished by a high percentage of materials with the  $Fm\bar{3}m$  space group (60%); this percentage is over twice that for the entire test set (26%). The low interpretability of descriptions associated with cubic crystals may require further analysis. In the case of sufficiency, the score distributions of all explainers have the main peak located at zero. Entities with a preferable score value (<0.1) hold a share of 97% (LIME), 95% (SHAP), 86% (SM), and 69% (IG). Low sufficiency means that only a tiny set of tokens actually affects model output. To sum up, LIME and SHAP provide rationales that are (often) comprehensible and (nearly always) sufficient to explain how a classifier in question makes a decision. There are additional factors to take into account for a holistic evaluation of language models beyond the faithfulness of XAI techniques. The high resource requirements of explainability evaluations make computation efficiency especially relevant. For instance, computation time for inference per test example is 12 ms, which is substantially less than the average times of 1.7 and 10.4 s taken by explainers SHAP and LIME, respectively. In the following analysis, the explainer with the highest rate of well-explained descriptions (LIME) is explored; other options should be considered in the scenarios where computation efficiency is a top priority.

Two classes (metals vs. nonmetals) are considered separately. Fifty most numerous unique tokens in both cases are finally examined (Figure 4); the corresponding data for other classification tasks are presented in Figures S1–S4. The visualized words can be formally categorized into two groups. Namely, the former consists of chemical-element symbols and associated subtokens ([##Bi], [##Ga], [##Sb], and [##Te]), while the latter contains the tokens accompanied by crystal structure. Although the MatBERT tokenizer was prepared on the in-domain corpus, it still cannot properly handle a minor part of chemical formulas. For instance,  $Rb_2IrF_6Hg$  decomposes into the following string of subtokens: [Rb][##2][##Ir][##F6][##Hg]. Therefore, the high priority of the [##F6] subtoken may indicate not only an influence of fluorine on a target property but also the importance of specific stoichiometry. The same conclusion is also true for similar subtokens depicted in Figure 4: [##F3], [##O3], [##Te2], and others. Unambiguous identification of stoichiometry's impact would require further developments in tokenization of chemical formulas. The MatBERT tokenizer does not ideally parse space group symbols either. For this reason, tokens [Pm] and [3m] are present due to incorrect processing of some space groups that inherit inversion symmetry:  $R\bar{3}m$ ,  $Pm\bar{3}m$ ,  $Fm\bar{3}m$ , and others. Token [mm] originates from splitting of space group symbols by a slash, e.g.,  $P4/mmm$  is split into [P4][/][mm][##m]. Here again, we cannot differentiate the importance of a specific space group and that of its inherent symmetry elements. Next, numbers as tokens ([one], [four], [six], [eight], and [12]) are related to the number of nearest neighbors of a described site. Then, the set of tokens, including [coplanar], [cubic], [cub][##octa][##hedral], [tetrahedral], [octahedral], [pyramidal], [trigonal], [water][\_] [like], and [hexagonal], serves to describe a coordination environment. It partially overlaps with another set, which contains crystal systems: [cubic], [trigonal], [hexagonal], [tetragonal], and [orthorhombic]. Tokens [Fluor][##ite], [Hal][##ite], and [Heusler] refer to the eponymous structural types. Finally, tokens [distorted] and [equivalent] help to characterize local (dis)order.



**Figure 3. Faithfulness metrics determined within *post hoc* local explainability techniques**

(A–H) The MatBERT model for band gap classification is examined. Each subplot contains a distribution of calculated metric values (comprehensiveness or sufficiency) and a respective cumulative curve. The results are presented for the following explainers: (A and E) Shapley additive explanations (SHAPs), (B and F) local interpretable model-agnostic explanations (LIMEs), (C and G) SMs, and (D and H) integrated gradients (IGs). The preferable directions for changing explainability measures are marked by arrows.

## DISCUSSION

Referring to a recent perspective,<sup>121</sup> the presented language-centric XAI framework interacts with two of three dimensions of AI-assisted scientific understanding. On the

Word clouds displayed in Figure 4 provide a bird’s-eye view of most general relations between a material’s features and its electronic structure (the presence of a band gap). MatBERT can reveal well-known patterns, such as the abundance of tetrahedral structures among semiconductors (nonmetals in our terminology) and dominance of intermetallic compounds as metals. To gain more insights into the plausibility of the presented language-centric approach, we directly compare most influential tokens extracted by a faithful explainer (LIME) with rationales proposed by a domain expert. It should be mentioned that the coauthor who highlighted meaningful tokens did not participate in ML model reasoning analysis. In this way, we sought to avoid a bias in human decision making. The domain expert operates under a discrete regime assigning one of two scores to each token (Figure 5): insignificant (0) or important (1). Then, his or her rationales are matched with continuous importance measures identified by the XAI method. Two plausibility metrics described in the section **experimental procedures** are calculated for a subset of test examples (287 entities, band gap classification): the token level F1 score and area under the precision-recall curve (AUPRC), which are equal to 0.33 and 0.32, respectively. Due to the absence of relevant XAI studies in the materials science field, we have to compare the obtained values with the values available in other scientific fields. In the Evaluating Rationales And Simple English Reasoning<sup>116</sup> (ERASER) benchmark study, seven datasets covering diverse document types (from reports of clinical trials to movie reviews) are analyzed to quantify interpretability of several language models. Best-performing models have a token-level F1 score and AUPRC in the range of 0.134–0.812 and 0.244–0.606, respectively. Thus, the consistency of our language model within domain expert reasoning is comparable to the previously obtained ones. We hope the present study will stimulate the creation of interpretability-aware benchmarks, resulting in an understanding of how one can reach high interpretability of ML algorithms in materials research. To facilitate such efforts, open access to the first expert-annotated corpus of materials descriptions for band gap predictions is provided.

one hand, scientists are supposed to generalize insights gained from a “computational microscope” without the need for complete computation. New data representations with inherent transparency can promote advances in the field. Indeed, crystal structure descriptions placed in the context of supervised ML help us to reduce model reasoning to the concepts that are well known to materials scientists. On the other hand, the suggested approach falls into the second category of AI contributions. We explicate internal machinery of language models within XAI techniques, and this approach will enable researchers to obtain unexpected results by inspecting originally black-boxed algorithms in the future. In the aforementioned dimensions, AI serves as a consultant for human scientists. On the contrary, the algorithms belonging to the third dimension<sup>121</sup> are thought to be independent agents of understanding capable of translating their vision. So far, there are no algorithms that undoubtedly fall into this category. Nonetheless, taking into account recent advances in large language models such as ChatGPT,<sup>122</sup> ultra-strong XAI approaches<sup>123</sup> have a bright future.


To sum up, here we present a language-centric framework aimed at accurately predicting materials properties and at providing clear explanations of the corresponding rationales simultaneously. State-of-the-art performance is grounded in the incorporation of domain knowledge into advanced transformer models through pretraining on a large corpus of papers. On the other hand, human-readable text-based descriptions as materials representation allow us to compare model reasoning and expert decisions directly. The proposed approach offers an alternative to opaque ML models, which are omnipresent in materials informatics at present. At the concept level, our intention was to dispel the popular belief that AI techniques are black boxes unable to stimulate new insights into structure-property relationships.

## EXPERIMENTAL PROCEDURES

### Resource availability

#### Lead contact

Further information and requests for resources should be directed to and will be fulfilled by the lead contact, Vadim Korolev ([korolev@colloid.chem.msu.ru](mailto:korolev@colloid.chem.msu.ru)).



**Figure 4. Most influential tokens in band gap classification provided by the *post hoc* local explainer**

Word clouds contain tokens that have the greatest impact on the MatBERT model decision to classify a material as a metal (left) or a nonmetal (right). The font size reflects the amplitude of the influence, whereas color differentiation helps to distinguish adjacent tokens.

### Materials availability

This study did not generate new unique reagents.

### Data and code availability

The expert-annotated dataset for band gap classification is available at Zenodo: <https://doi.org/10.5281/zenodo.7750192>. The trained MatBERT model for band gap classification is available at Hugging Face Hub: <https://huggingface.co/korolewadim/matbert-bandgap> and Zenodo: <https://doi.org/10.5281/zenodo.7992527>. The source code and data accompanying this work are publicly available at Zenodo: <https://doi.org/10.5281/zenodo.7992558>.

### Datasets

The JARVIS-DFT dataset, a part of JARVIS,<sup>97</sup> was the main source of data in the study. The vdW-DF-OptB88 van der Waals functional<sup>124</sup> was used to calculate most materials properties; the Tran-Blaha modified Becke-Johnson functional<sup>125</sup> was selected to reproduce better band gaps<sup>126</sup> and frequency-dependent dielectric functions and hence SLME.<sup>127</sup> The following classification tasks were analyzed: energy above the convex hull (55,350 entities, a threshold of 0.1 eV/atom), the magnetic moment (52,205, 0.05  $\mu$ B), a band gap (18,167, 0.01 eV), SLME (9,063, 10%), and topological spin-orbit spillage (11,376, 0.1).

### Model training

k-Fold cross-validation (10-fold) was performed to estimate model performance. In the case of attention-based neural networks trained on composition and graph neural networks, one-ninth of training data were retained as a validation set for early stopping, thereby ensuring a training-validation-test ratio of 80:10:10. The data from the validation set were not used in the training of all other models. We employed identical data splits for each prediction algorithm, considering a specific dataset. The analyzed classification metrics (accuracy, the F1 score, and MCC) were averaged over cross-validation subsets, while standard deviation was regarded as a measure of prediction uncertainty. All the deep learning models were developed within the PyTorch framework.<sup>128</sup>

### Random forests on force-field-inspired descriptors

We trained random forest<sup>105</sup> on CFIDs,<sup>106</sup> including chemical, cell size, radial charge, and distribution (radial, angular, dihedral, and nearest-neighbor) features. The scikit-learn implementation<sup>129</sup> of the algorithm was chosen with default hyperparameter values. CFIDs were extracted using the JARVIS-Tools library.<sup>97</sup>

### Attention-based neural networks on composition

A neural network referred to as the Roost framework<sup>107</sup> was trained on materials compositions represented as dense weighted graphs between elements. Node representations were updated through message passage by the soft-attention mechanism. Fixed-length materials representations generated via a soft-attention pooling operation were then passed as input to a feedforward network that finally generated an endpoint value. The AdamW optimizer<sup>130</sup> with parameters  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ , and a learning rate of  $10^{-3}$  was employed. The model was trained for 1,000 epochs with an early stopping at 100. The batch size was set to 128.

### Graph neural networks

We trained ALIGNN,<sup>18</sup> which was intended to capture many-body interactions explicitly. The ALIGNN architecture performed a series of edge-gated graph

convolutions on an atomistic and corresponding line graph. The resulting atom representations were reduced by an average pooling operation and transferred to a fully connected network to predict a target property. The AdamW optimizer<sup>130</sup> with parameters  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ , and a learning rate of  $10^{-3}$  was employed. The model was trained for 1,000 epochs with an early stopping at 100. The batch size was set to 64. The original ALIGNN implementation was used that heavily relies on the Deep Graph Library.<sup>131</sup>

### Transformer language models

The human-readable descriptions of crystal structure were generated by means of the Robocrystallographer library.<sup>101</sup> The text information on the local (coordination number and geometry), semilocal (polyhedral connectivity and tilts angles), and global (mineral type and crystal symmetry) environments was represented as a sequence of tokens. The BERT<sup>103</sup> model was chosen as a basic architecture. Taking into account weights' initialization and tokenization procedures, three models were trained for each downstream task: a randomly initialized BERT model using the original tokenizer, a randomly initialized BERT model using the MatBERT tokenizer, and the MatBERT model<sup>93</sup> using the MatBERT tokenizer. Both case-sensitive tokenizers were based on the WordPiece algorithm.<sup>132</sup> The AdamW optimizer<sup>130</sup> with parameters  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ , and a learning rate of  $3 \times 10^{-4}$  was utilized. The model was trained over 10 epochs with a batch size of 16. The HuggingFace Transformers library<sup>133</sup> was extensively used to assess pretrained models and to fine-tune them.

JVASP-2133 is classified as nonmetal  
(tokens that **support** and **surpress** prediction)

### AI rationales

[CLS] Hg **##O** crystallizes in the **orthorhombic** Pn **##ma** space group . Hg ( 1 ) is bonded in a linear geometry to two equivalent **O** ( 1 ) atoms . There is one **shorter** ( 2 . 06 Å ) and one **longer** ( 2 . 07 Å ) Hg ( 1 ) - **O** ( 1 ) bond length . **O** ( 1 ) is bonded in a **water** - like geometry to two equivalent Hg ( 1 ) atoms . [SEP]

### human rationales

[CLS] Hg **##O** crystallizes in the **orthorhombic** Pn **##ma** space group . Hg ( 1 ) is bonded in a **linear** geometry to two equivalent **O** ( 1 ) atoms . There is one **shorter** ( 2 . 06 Å ) and one **longer** ( 2 . 07 Å ) Hg ( 1 ) - O ( 1 ) bond length . O ( 1 ) is bonded in a **water** - like geometry to **two** equivalent Hg ( 1 ) atoms . [SEP]

**Figure 5. Token importance levels determined by a *post hoc* local explainer and a domain expert**

Distinct words are colored in accordance with their impact on model output; token intensity denotes the amplitude, whereas the color (red vs. blue) means a contribution to the predicted class (negative vs. positive).



## Model explanation

### Explanation algorithms

We took advantage of four XAI techniques. First, SMs<sup>117</sup> of the predicted class were generated. The method was originally formulated for image-specific class saliency visualization; here, elements of SM, i.e., feature importance levels, were extracted at the token level as derivatives of predicted class probability with respect to the corresponding token embedding. Second, IGs<sup>118</sup> were defined as path integrals of the gradients along the straight-line path from the baseline (the padding token) to the considered token. Both SM and IG were calculated using the Captum package.<sup>134</sup> Third, the LIME<sup>119</sup> approach was implemented to obtain token level importance scores. The predictors in question were approximated by a transparent algorithm: Ridge regression. Then, the surrogate model was optimized in such a way as to ensure both interpretability and local fidelity. We employed the original implementation of the algorithm for this purpose. Fourth, Shapley values<sup>135</sup> from game theory were assigned in order to quantify tokens' contributions to the model outcome. To be precise, the extended version of Shapley values, also called Owen values,<sup>136</sup> was computed to capture preferable input feature coalitions. Partition masking, as implemented in the SHAP package,<sup>54</sup> was applied for this purpose.

### Evaluation metrics

Faithfulness of ML predictors was assessed via two metrics.<sup>116</sup> Starting with an original sequence of tokens  $x_i$ , we constructed its contrast example by removing subset of tokens  $r_i$ . Comprehensiveness is defined as a difference between probability assigned by model  $m$  to initial sequence of tokens  $x_i$  and probability derived by the same algorithm from sequence with removed rationales  $x_i \setminus r_i$ :

$$\text{comprehensiveness} = m(x_i) - m(x_i \setminus r_i)$$

Sufficiency is oppositely defined as a difference between probability assigned by model  $m$  to initial sequence of tokens  $x_i$  and probability derived by the same algorithm from sequence of removed rationales  $r_i$ :

$$\text{sufficiency} = m(x_i) - m(r_i)$$

Both metrics were calculated for predicted class  $j$  (i.e., the class with the highest probability  $m(x_i)_j$ ). The arbitrariness of the choice of subset  $r_i$  was overcome as follows. We calculated faithfulness measures assuming subsets of rationales  $r_{ik}$  that included  $k$  percent of most important tokens identified by an explainer of interest. Then, an aggregate metric referred to as area over the perturbation curve<sup>116</sup> (AOPC) was calculated as

$$\text{AOPC} = \frac{1}{|\mathfrak{B}|+1} \left( \sum_{k=0}^{|\mathfrak{B}|} m(x_i)_j - m(x_i \setminus r_{ik})_j \right)$$

The set of percentiles  $\mathfrak{B}$  is  $\{10\%, 20\%, \dots, 100\%\}$ . Throughout the main text, by comprehensiveness and sufficiency we mean the corresponding AOPC values.

Plausibility was estimated for discrete and soft explanations.<sup>116</sup> For each example, the subset of tokens selected by a domain expert was compared with the subset of rationales that included  $k$  most influential tokens according to explainer ranking. The value of  $k$  is set to the average rationale length proposed by a human (10). The corresponding token level F1 score was regarded as a plausibility measure. In addition, we estimated AUPRC to take into account tokens' ranking.

All the explainability-related calculations were carried out within the ferret package.<sup>104</sup>

## SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.patter.2023.100803>.

## AUTHOR CONTRIBUTIONS

Conceptualization, methodology, software, writing – original draft, supervision, V.K.; investigation, writing – review and editing, V.K. and P.P.

## DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: March 28, 2023

Revised: June 6, 2023

Accepted: July 4, 2023

Published: August 2, 2023

## REFERENCES

- von Lilienfeld, O.A. (2020). Introducing machine learning: science and technology. *Mach. Learn. Sci. Technol.* 1, 010201. <https://doi.org/10.1088/2632-2153/ab6d5d>.
- Meredig, B., Agrawal, A., Kirklín, S., Saal, J.E., Doak, J.W., Thompson, A., Zhang, K., Choudhary, A., and Wolverton, C. (2014). Combinatorial screening for new materials in unconstrained composition space with machine learning. *Phys. Rev. B* 89, 094104. <https://doi.org/10.1103/PhysRevB.89.094104>.
- Faber, F.A., Lindmaa, A., Von Lilienfeld, O.A., and Armiento, R. (2016). Machine learning energies of 2 million elpasolite (A B C 2 D 6) crystals. *Phys. Rev. Lett.* 117, 135502. <https://doi.org/10.1103/PhysRevLett.117.135502>.
- Seko, A., Hayashi, H., Nakayama, K., Takahashi, A., and Tanaka, I. (2017). Representation of compounds for machine-learning prediction of physical properties. *Phys. Rev. B* 95, 144110. <https://doi.org/10.1103/PhysRevB.95.144110>.
- Bartel, C.J., Millican, S.L., Deml, A.M., Rumptz, J.R., Tumas, W., Weimer, A.W., Lany, S., Stevanović, V., Musgrave, C.B., and Holder, A.M. (2018). Physical descriptor for the Gibbs energy of inorganic crystalline solids and temperature-dependent materials chemistry. *Nat. Commun.* 9, 4168. <https://doi.org/10.1038/s41467-018-06682-4>.
- Noh, J., Kim, J., Stein, H.S., Sanchez-Lengeling, B., Gregoire, J.M., Aspuru-Guzik, A., and Jung, Y. (2019). Inverse design of solid-state materials via a continuous representation. *Matter* 1, 1370–1384. <https://doi.org/10.1016/j.matt.2019.08.017>.
- Korolev, V., Mitrofanov, A., Eliseev, A., and Tkachenko, V. (2020). Machine-learning-assisted search for functional materials over extended chemical space. *Mater. Horiz.* 7, 2710–2718. <https://doi.org/10.1039/D0MH00881H>.
- Yao, Z., Sánchez-Lengeling, B., Bobbitt, N.S., Bucior, B.J., Kumar, S.G.H., Collins, S.P., Burns, T., Woo, T.K., Farha, O.K., Snurr, R.Q., and Aspuru-Guzik, A. (2021). Inverse design of nanoporous crystalline reticular materials with deep generative models. *Nat. Mach. Intell.* 3, 76–86. <https://doi.org/10.1038/s42256-020-00271-1>.
- Ren, Z., Tian, S.I.P., Noh, J., Oviedo, F., Xing, G., Li, J., Liang, Q., Zhu, R., Aberle, A.G., Sun, S., et al. (2022). An invertible crystallographic representation for general inverse design of inorganic crystals with targeted properties. *Matter* 5, 314–335. <https://doi.org/10.1016/j.matt.2021.11.032>.
- Chen, C., and Ong, S.P. (2022). A universal graph deep learning interatomic potential for the periodic table. *Nat. Comput. Sci.* 2, 718–728. <https://doi.org/10.1038/s43588-022-00349-3>.
- Sauceda, H.E., Gálvez-González, L.E., Chmiela, S., Paz-Borbón, L.O., Müller, K.R., and Tkatchenko, A. (2022). BIGDML—Towards accurate quantum machine learning force fields for materials. *Nat. Commun.* 13, 3733. <https://doi.org/10.1038/s41467-022-31093-x>.
- Choudhary, K., DeCost, B., Major, L., Butler, K., Thiyagalingam, J., and Tavazza, F. (2023). Unified Graph Neural Network Force-field for the Periodic Table for Solids. *Dig. Dis.* 2, 346–355. <https://doi.org/10.1039/D2DD00096B>.
- Ryan, K., Lengyel, J., and Shatruk, M. (2018). Crystal structure prediction via deep learning. *J. Am. Chem. Soc.* 140, 10158–10168. <https://doi.org/10.1021/jacs.8b03913>.

14. Podryabinkin, E.V., Tikhonov, E.V., Shapeev, A.V., and Oganov, A.R. (2019). Accelerating crystal structure prediction by machine-learning interatomic potentials with active learning. *Phys. Rev. B* 99, 064114. <https://doi.org/10.1103/PhysRevB.99.064114>.
15. Liang, H., Stanev, V., Kusne, A.G., and Takeuchi, I. (2020). CRYSPNet: Crystal structure predictions via neural networks. *Phys. Rev. Mater.* 4, 123802. <https://doi.org/10.1103/PhysRevMaterials.4.123802>.
16. Reiser, P., Neubert, M., Eberhard, A., Torresi, L., Zhou, C., Shao, C., Metni, H., van Hoesel, C., Schopmans, H., Sommer, T., and Friederich, P. (2022). Graph neural networks for materials science and chemistry. *Commun. Mater.* 3, 93. <https://doi.org/10.1038/s43246-022-00315-6>.
17. Chen, C., Ye, W., Zuo, Y., Zheng, C., and Ong, S.P. (2019). Graph networks as a universal machine learning framework for molecules and crystals. *Chem. Mater.* 31, 3564–3572. <https://doi.org/10.1021/acs.chemmater.9b01294>.
18. Choudhary, K., and DeCost, B. (2021). Atomistic line graph neural network for improved materials property predictions. *npj Comput. Mater.* 7, 185. <https://doi.org/10.1038/s41524-021-00650-1>.
19. Xie, T., and Grossman, J.C. (2018). Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties. *Phys. Rev. Lett.* 120, 145301. <https://doi.org/10.1103/PhysRevLett.120.145301>.
20. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. (2020). Language models are few-shot learners. *Adv. Neural Inf. Process. Syst.* 33, 1877–1901.
21. Fedus, W., Zoph, B., and Shazeer, N. (2022). Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *J. Mach. Learn. Res.* 23, 1–40.
22. Holm, E.A. (2019). In defense of the black box. *Science* 364, 26–27. <https://doi.org/10.1126/science.aax0162>.
23. Doshi-Velez, F., and Kim, B. (2017). Towards a rigorous science of interpretable machine learning. Preprint at arXiv. <https://doi.org/10.48550/arXiv.1702.08608>.
24. Meredig, B. (2019). Five high-impact research areas in machine learning for materials science. *Chem. Mater.* 31, 9579–9581. <https://doi.org/10.1021/acs.chemmater.9b04078>.
25. Gunning, D., Stefik, M., Choi, J., Miller, T., Stumpf, S., and Yang, G.-Z. (2019). XAI—Explainable artificial intelligence. *Sci. Robot.* 4, eaay7120. <https://doi.org/10.1126/scirobotics.aay7120>.
26. Gunning, D., and Aha, D. (2019). DARPA’s explainable artificial intelligence (XAI) program. *AI Mag.* 40, 44–58. <https://doi.org/10.1609/aimag.v40i2.2850>.
27. Adadi, A., and Berrada, M. (2018). Peeking inside the black-box: a survey on explainable artificial intelligence (XAI). *IEEE Access* 6, 52138–52160. <https://doi.org/10.1109/ACCESS.2018.2870052>.
28. Barredo Arrieta, A., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., García, S., Gil-López, S., Molina, D., Benjamins, R., et al. (2020). Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Inf. Fusion* 58, 82–115. <https://doi.org/10.1016/j.inffus.2019.12.012>.
29. Das, A., and Rad, P. (2020). Opportunities and challenges in explainable artificial intelligence (xai): A survey. Preprint at arXiv. <https://doi.org/10.48550/arXiv.2006.11371>.
30. Esterhuizen, J.A., Goldsmith, B.R., and Linic, S. (2022). Interpretable machine learning for knowledge generation in heterogeneous catalysis. *Nat. Catal.* 5, 175–184. <https://doi.org/10.1038/s41929-022-00744-z>.
31. Zhong, X., Gallagher, B., Liu, S., Kailkhura, B., Hiszpanski, A., and Han, T.Y.-J. (2022). Explainable machine learning in materials science. *npj Comput. Mater.* 8, 204. <https://doi.org/10.1038/s41524-022-00884-7>.
32. Oviedo, F., Ferrer, J.L., Buonassisi, T., and Butler, K.T. (2022). Interpretable and explainable machine learning for materials science and chemistry. *Acc. Mater. Res.* 3, 597–607. <https://doi.org/10.1021/ac-countsmr.1c00244>.
33. Omidvar, N., Pillai, H.S., Wang, S.-H., Mou, T., Wang, S., Athawale, A., Achenie, L.E.K., and Xin, H. (2021). Interpretable machine learning of chemical bonding at solid surfaces. *J. Phys. Chem. Lett.* 12, 11476–11487. <https://doi.org/10.1021/acs.jpcclett.1c03291>.
34. Lou, Y., Caruana, R., and Gehrke, J. (2012). Intelligible models for classification and regression. In Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining, pp. 150–158. <https://doi.org/10.1145/2339530.2339556>.
35. Lou, Y., Caruana, R., Gehrke, J., and Hooker, G. (2013). Accurate intelligible models with pairwise interactions. In Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining, pp. 623–631. <https://doi.org/10.1145/2487575.2487579>.
36. Xin, H., Holewinski, A., and Linic, S. (2012). Predictive structure–reactivity models for rapid screening of Pt-based multimetallic electrocatalysts for the oxygen reduction reaction. *ACS Catal.* 2, 12–16. <https://doi.org/10.1021/cs200462f>.
37. Allen, A.E.A., and Tkatchenko, A. (2022). Machine learning of material properties: Predictive and interpretable multilinear models. *Sci. Adv.* 8, eabm7185. <https://doi.org/10.1126/sciadv.abm7185>.
38. Esterhuizen, J.A., Goldsmith, B.R., and Linic, S. (2020). Theory-guided machine learning finds geometric structure–property relationships for chemisorption on subsurface alloys. *Chem* 6, 3100–3117. <https://doi.org/10.1016/j.chempr.2020.09.001>.
39. Ishikawa, A., Sodeyama, K., Igarashi, Y., Nakayama, T., Tateyama, Y., and Okada, M. (2019). Machine learning prediction of coordination energies for alkali group elements in battery electrolyte solvents. *Phys. Chem. Chem. Phys.* 21, 26399–26405. <https://doi.org/10.1039/C9CP03679B>.
40. Breiman, L. (2017). *Classification and Regression Trees* (Routledge).
41. Carrete, J., Mingo, N., Wang, S., and Curtarolo, S. (2014). Nanograined Half-Heusler Semiconductors as Advanced Thermoelectrics: An Ab Initio High-Throughput Statistical Study. *Adv. Funct. Mater.* 24, 7427–7432. <https://doi.org/10.1002/adfm.201401201>.
42. Fernandez, M., and Barnard, A.S. (2016). Geometrical properties can predict CO<sub>2</sub> and N<sub>2</sub> adsorption performance of metal–organic frameworks (MOFs) at low pressure. *ACS Comb. Sci.* 18, 243–252. <https://doi.org/10.1021/acscombsci.5b00188>.
43. Wu, Y., Duan, H., and Xi, H. (2020). Machine learning-driven insights into defects of zirconium metal–organic frameworks for enhanced ethane–ethylene separation. *Chem. Mater.* 32, 2986–2997. <https://doi.org/10.1021/acs.chemmater.9b05322>.
44. Wang, Y., Wagner, N., and Rondinelli, J.M. (2019). Symbolic regression in materials science. *MRS Commun.* 9, 793–805. <https://doi.org/10.1557/mrc.2019.85>.
45. Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *J. R. Stat. Soc. Ser. B Methodol.* 58, 267–288. <https://doi.org/10.1111/j.2517-6161.1996.tb02080.x>.
46. Ouyang, R., Curtarolo, S., Ahmetcik, E., Scheffler, M., and Ghiringhelli, L.M. (2018). SISSO: A compressed-sensing method for identifying the best low-dimensional descriptor in an immensity of offered candidates. *Phys. Rev. Mater.* 2, 083802. <https://doi.org/10.1103/PhysRevMaterials.2.083802>.
47. Ghiringhelli, L.M., Vybiral, J., Levchenko, S.V., Draxl, C., and Scheffler, M. (2015). Big data of materials science: critical role of the descriptor. *Phys. Rev. Lett.* 114, 105503. <https://doi.org/10.1103/PhysRevLett.114.105503>.
48. Andersen, M., Levchenko, S.V., Scheffler, M., and Reuter, K. (2019). Beyond scaling relations for the description of catalytic materials. *ACS Catal.* 9, 2752–2759. <https://doi.org/10.1021/acscatal.8b04478>.
49. Bartel, C.J., Sutton, C., Goldsmith, B.R., Ouyang, R., Musgrave, C.B., Ghiringhelli, L.M., and Scheffler, M. (2019). New tolerance factor to predict the stability of perovskite oxides and halides. *Sci. Adv.* 5, eaav0693. <https://doi.org/10.1126/sciadv.aav0693>.
50. Weng, B., Song, Z., Zhu, R., Yan, Q., Sun, Q., Grice, C.G., Yan, Y., and Yin, W.-J. (2020). Simple descriptor derived from symbolic regression accelerating the discovery of new perovskite catalysts. *Nat. Commun.* 11, 3513. <https://doi.org/10.1038/s41467-020-17263-9>.

51. Singstock, N.R., Ortiz-Rodríguez, J.C., Perryman, J.T., Sutton, C., Velázquez, J.M., and Musgrave, C.B. (2021). Machine learning guided synthesis of multinary chevre phase chalcogenides. *J. Am. Chem. Soc.* *143*, 9113–9122. <https://doi.org/10.1021/jacs.1c02971>.
52. Cao, G., Ouyang, R., Ghiringhelli, L.M., Scheffler, M., Liu, H., Carbogno, C., and Zhang, Z. (2020). Artificial intelligence for high-throughput discovery of topological insulators: The example of alloyed tetradymites. *Phys. Rev. Mater.* *4*, 034204. <https://doi.org/10.1103/PhysRevMaterials.4.034204>.
53. Xu, W., Andersen, M., and Reuter, K. (2020). Data-driven descriptor engineering and refined scaling relations for predicting transition metal oxide reactivity. *ACS Catal.* *11*, 734–742. <https://doi.org/10.1021/acscatal.0c04170>.
54. Lundberg, S.M., and Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Adv. Neural Inf. Process. Syst.* *30*.
55. Liang, J., and Zhu, X. (2019). Phillips-inspired machine learning for band gap and exciton binding energy prediction. *J. Phys. Chem. Lett.* *10*, 5640–5646. <https://doi.org/10.1021/acs.jpcclett.9b02232>.
56. Korolev, V.V., Mitrofanov, A., Marchenko, E.I., Eremin, N.N., Tkachenko, V., and Kalmykov, S.N. (2020). Transferable and extensible machine learning-derived atomic charges for modeling hybrid nanoporous materials. *Chem. Mater.* *32*, 7822–7831. <https://doi.org/10.1021/acs.chemmater.0c02468>.
57. Jablonka, K.M., Moosavi, S.M., Asgari, M., Ireland, C., Patiny, L., and Smit, B. (2020). A data-driven perspective on the colours of metal-organic frameworks. *Chem. Sci.* *12*, 3587–3598. <https://doi.org/10.1039/D0SC05337F>.
58. Georgescu, A.B., Ren, P., Toland, A.R., Zhang, S., Miller, K.D., Apley, D.W., Olivetti, E.A., Wagner, N., and Rondinelli, J.M. (2021). Database, Features, and Machine Learning Model to Identify Thermally Driven Metal-Insulator Transition Compounds. *Chem. Mater.* *33*, 5591–5605. <https://doi.org/10.1021/acs.chemmater.1c00905>.
59. Zhang, S., Lu, T., Xu, P., Tao, Q., Li, M., and Lu, W. (2021). Predicting the Formability of Hybrid Organic-Inorganic Perovskites via an Interpretable Machine Learning Strategy. *J. Phys. Chem. Lett.* *12*, 7423–7430. <https://doi.org/10.1021/acs.jpcclett.1c01939>.
60. Marchenko, E.I., Korolev, V.V., Fateev, S.A., Mitrofanov, A., Eremin, N.N., Goodilin, E.A., and Tarasov, A.B. (2021). Relationships between distortions of inorganic framework and band gap of layered hybrid halide perovskites. *Chem. Mater.* *33*, 7518–7526. <https://doi.org/10.1021/acs.chemmater.1c02467>.
61. Korolev, V.V., Nevolin, Y.M., Manz, T.A., and Protsenko, P.V. (2021). Parametrization of Nonbonded Force Field Terms for Metal-Organic Frameworks Using Machine Learning Approach. *J. Chem. Inf. Model.* *61*, 5774–5784. <https://doi.org/10.1021/acs.jcim.1c01124>.
62. Anker, A.S., Kjær, E.T.S., Juelsholt, M., Christiansen, T.L., Skjærvø, S.L., Jørgensen, M.R.V., Kantor, I., Sørensen, D.R., Billinge, S.J.L., Selvan, R., and Jensen, K.M.Ø. (2022). Extracting structural motifs from pair distribution function data of nanostructures using explainable machine learning. *npj Comput. Mater.* *8*, 213. <https://doi.org/10.1038/s41524-022-00896-3>.
63. Lu, S., Zhou, Q., Guo, Y., and Wang, J. (2022). On-the-fly interpretable machine learning for rapid discovery of two-dimensional ferromagnets with high Curie temperature. *Chem* *8*, 769–783. <https://doi.org/10.1016/j.chempr.2021.11.009>.
64. Wu, Y., Lu, S., Ju, M.-G., Zhou, Q., and Wang, J. (2021). Accelerated design of promising mixed lead-free double halide organic-inorganic perovskites for photovoltaics using machine learning. *Nanoscale* *13*, 12250–12259. <https://doi.org/10.1039/D1NR01117K>.
65. Kronberg, R., Lappalainen, H., and Laasonen, K. (2021). Hydrogen adsorption on defective nitrogen-doped carbon nanotubes explained via machine learning augmented DFT calculations and game-theoretic feature attributions. *J. Phys. Chem. C* *125*, 15918–15933. <https://doi.org/10.1021/acs.jpcc.1c03858>.
66. Lu, T., Li, H., Li, M., Wang, S., and Lu, W. (2022). Predicting Experimental Formability of Hybrid Organic-Inorganic Perovskites via Imbalanced Learning. *J. Phys. Chem. Lett.* *13*, 3032–3038. <https://doi.org/10.1021/acs.jpcclett.2c00603>.
67. Lipton, Z.C. (2018). The mythos of model interpretability: In machine learning, the concept of interpretability is both important and slippery. *Queue* *16*, 31–57. <https://doi.org/10.1145/3236386.3241340>.
68. Musil, F., Grisafi, A., Bartók, A.P., Ortner, C., Csányi, G., and Ceriotti, M. (2021). Physics-inspired structural representations for molecules and materials. *Chem. Rev.* *121*, 9759–9815. <https://doi.org/10.1021/acs.chemrev.1c00021>.
69. Battaglia, P.W., Hamrick, J.B., Bapst, V., Sanchez-Gonzalez, A., Zambaldi, V., Malinowski, M., Tacchetti, A., Raposo, D., Santoro, A., Faulkner, R., et al. (2018). Relational inductive biases, deep learning, and graph networks. Preprint at arXiv. <https://doi.org/10.48550/arXiv.1806.01261>.
70. Bronstein, M.M., Bruna, J., Cohen, T., and Velicković, P. (2021). Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. Preprint at arXiv. <https://doi.org/10.48550/arXiv.2104.13478>.
71. Fedorov, A.V., and Shamanaev, I.V. (2017). Crystal structure representation for neural networks using topological approach. *Mol. Inform.* *36*, 1600162. <https://doi.org/10.1002/minf.201600162>.
72. Korolev, V., Mitrofanov, A., Korotcov, A., and Tkachenko, V. (2020). Graph convolutional neural networks as “general-purpose” property predictors: the universality and limits of applicability. *J. Chem. Inf. Model.* *60*, 22–28. <https://doi.org/10.1021/acs.jcim.9b00587>.
73. Park, C.W., and Wolverton, C. (2020). Developing an improved crystal graph convolutional neural network framework for accelerated materials discovery. *Phys. Rev. Mater.* *4*, 063801. <https://doi.org/10.1103/PhysRevMaterials.4.063801>.
74. Cheng, J., Zhang, C., and Dong, L. (2021). A geometric-information-enhanced crystal graph network for predicting properties of materials. *Commun. Mater.* *2*, 92. <https://doi.org/10.1038/s43246-021-00194-3>.
75. Omeel, S.S., Louis, S.-Y., Fu, N., Wei, L., Dey, S., Dong, R., Li, Q., and Hu, J. (2022). Scalable deeper graph neural networks for high-performance materials property prediction. *Patterns* *3*, 100491. <https://doi.org/10.1016/j.patter.2022.100491>.
76. Fung, V., Zhang, J., Juarez, E., and Sumpter, B.G. (2021). Benchmarking graph neural networks for materials chemistry. *npj Comput. Mater.* *7*, 84. <https://doi.org/10.1038/s41524-021-00554-0>.
77. Ying, R., Bourgeois, D., You, J., Zitnik, M., and Leskovec, J. (2019). Gnnexplainer: Generating explanations for graph neural networks. *Adv. Neural Inf. Process. Syst.* *32*, 9240–9251.
78. Noutahi, E., Beaini, D., Horwood, J., Giguère, S., and Tossou, P. (2019). Towards interpretable sparse graph representation learning with laplacian pooling. Preprint at arXiv. <https://doi.org/10.48550/arXiv.1905.11577>.
79. Luo, D., Cheng, W., Xu, D., Yu, W., Zong, B., Chen, H., and Zhang, X. (2020). Parameterized explainer for graph neural network. *Adv. Neural Inf. Process. Syst.* *33*, 19620–19631.
80. Pope, P.E., Kolouri, S., Rostami, M., Martin, C.E., and Hoffmann, H. (2019). Explainability methods for graph convolutional neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10772–10781.
81. Sun, J., Lapuschkin, S., Samek, W., Zhao, Y., Cheung, N.-M., and Binder, A. (2021). Explanation-guided training for cross-domain few-shot classification. In *2020 25th International Conference on Pattern Recognition (ICPR) (IEEE)*, pp. 7609–7616. <https://doi.org/10.1109/ICPR48806.2021.9412941>.
82. Huang, Q., Yamada, M., Tian, Y., Singh, D., and Chang, Y. (2022). Graphlime: Local interpretable model explanations for graph neural networks. *IEEE Trans. Knowl. Data Eng.* *1–6*. <https://doi.org/10.1109/TKDE.2022.3187455>.
83. Raza, A., Waqar, F., Sturluson, A., Simon, C., and Fern, X. (2020). Towards explainable message passing networks for predicting carbon dioxide adsorption in metal-organic frameworks. Preprint at arXiv. <https://doi.org/10.48550/arXiv.2012.03723>.
84. Hsu, T., Pham, T.A., Keilbart, N., Weitzner, S., Chapman, J., Xiao, P., Qiu, S.R., Chen, X., and Wood, B.C. (2022). Efficient and interpretable graph



- network representation for angle-dependent properties applied to optical spectroscopy. *npj Comput. Mater.* **8**, 151. <https://doi.org/10.1038/s41524-022-00841-4>.
85. Chen, P., Jiao, R., Liu, J., Liu, Y., and Lu, Y. (2022). Interpretable Graph Transformer Network for Predicting Adsorption Isotherms of Metal–Organic Frameworks. *J. Chem. Inf. Model.* **62**, 5446–5456. <https://doi.org/10.1021/acs.jcim.2c00876>.
  86. Montavon, G., Samek, W., and Müller, K.R. (2018). Methods for interpreting and understanding deep neural networks. *Digit. Signal Process.* **73**, 1–15. <https://doi.org/10.1016/j.dsp.2017.10.011>.
  87. Gilpin, L.H., Bau, D., Yuan, B.Z., Bajwa, A., Specter, M., and Kagal, L. (2018). Explaining Explanations: An Overview of Interpretability of Machine Learning (2018 IEEE 5th International Conference on data science and advanced analytics (DSAA) (IEEE)), pp. 80–89.
  88. Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., and Pedreschi, D. (2018). A survey of methods for explaining black box models. *ACM Comput. Surv.* **51**, 1–42. <https://doi.org/10.1145/3236009>.
  89. Schwartz, M.D. (2022). Should artificial intelligence be interpretable to humans? *Nat. Rev. Phys.* **4**, 741–742. <https://doi.org/10.1038/s42254-022-00538-z>.
  90. Hirschberg, J., and Manning, C.D. (2015). Advances in natural language processing. *Science* **349**, 261–266. <https://doi.org/10.1126/science.aaa8685>.
  91. Weston, L., Tshitoyan, V., Dagdelen, J., Kononova, O., Trewartha, A., Persson, K.A., Ceder, G., and Jain, A. (2019). Named entity recognition and normalization applied to large-scale information extraction from the materials science literature. *J. Chem. Inf. Model.* **59**, 3692–3702. <https://doi.org/10.1021/acs.jcim.9b00470>.
  92. Gupta, T., Zaki, M., Krishnan, N.M.A., and Mausam. (2022). MatSciBERT: A materials domain language model for text mining and information extraction. *npj Comput. Mater.* **8**, 102. <https://doi.org/10.1038/s41524-022-00784-w>.
  93. Trewartha, A., Walker, N., Huo, H., Lee, S., Cruse, K., Dagdelen, J., Dunn, A., Persson, K.A., Ceder, G., and Jain, A. (2022). Quantifying the advantage of domain-specific pre-training on named entity recognition tasks in materials science. *Patterns* **3**, 100488. <https://doi.org/10.1016/j.patter.2022.100488>.
  94. Huo, H., Rong, Z., Kononova, O., Sun, W., Botari, T., He, T., Tshitoyan, V., and Ceder, G. (2019). Semi-supervised machine-learning classification of materials synthesis procedures. *npj Comput. Mater.* **5**, 62. <https://doi.org/10.1038/s41524-019-0204-1>.
  95. Huang, S., and Cole, J.M. (2022). BatteryBERT: A Pretrained Language Model for Battery Database Enhancement. *J. Chem. Inf. Model.* **62**, 6365–6377. <https://doi.org/10.1021/acs.jcim.2c00035>.
  96. Haghightalari, M., Li, J., Heidar-Zadeh, F., Liu, Y., Guan, X., and Head-Gordon, T. (2020). Learning to make chemical predictions: the interplay of feature representation, data, and machine learning methods. *Chem* **6**, 1527–1542. <https://doi.org/10.1016/j.chempr.2020.05.014>.
  97. Choudhary, K., Garrity, K.F., Reid, A.C.E., DeCost, B., Biacchi, A.J., Hight Walker, A.R., Trautt, Z., Hatrick-Simpers, J., Kusne, A.G., Centrone, A., et al. (2020). The joint automated repository for various integrated simulations (JARVIS) for data-driven materials design. *npj Comput. Mater.* **6**, 173. <https://doi.org/10.1038/s41524-020-00440-1>.
  98. Curtarolo, S., Setyawan, W., Hart, G.L., Jahnatek, M., Chepulskii, R.V., Taylor, R.H., Wang, S., Xue, J., Yang, K., Levy, O., et al. (2012). AFLOW: An automatic framework for high-throughput materials discovery. *Comput. Mater. Sci.* **58**, 218–226. <https://doi.org/10.1016/j.com-matsci.2012.02.005>.
  99. Saal, J.E., Kirklín, S., Aykol, M., Meredig, B., and Wolverton, C. (2013). Materials design and discovery with high-throughput density functional theory: the open quantum materials database (OQMD). *Jom* **65**, 1501–1509. <https://doi.org/10.1007/s11837-013-0755-4>.
  100. Jain, A., Ong, S.P., Hautier, G., Chen, W., Richards, W.D., Dacek, S., Cholia, S., Gunter, D., Skinner, D., Ceder, G., and Persson, K.A. (2013). Commentary: The Materials Project: A materials genome approach to accelerating materials innovation. *Apl. Mater.* **1**, 011002. <https://doi.org/10.1063/1.4812323>.
  101. Ganose, A.M., and Jain, A. (2019). Robocrystallographer: automated crystal structure text descriptions and analysis. *MRS Commun.* **9**, 874–881. <https://doi.org/10.1557/mrc.2019.94>.
  102. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., and Polosukhin, I. (2017). Attention is all you need. *Adv. Neural Inf. Process. Syst.* **30**.
  103. Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. Preprint at arXiv. <https://doi.org/10.48550/arXiv.1810.04805>.
  104. Attanasio, G., Pastor, E., Di Bonaventura, C., and Nozza, D. (2022). ferret: a Framework for Benchmarking Explainers on Transformers. Preprint at arXiv. <https://doi.org/10.48550/arXiv.2208.01575>.
  105. Breiman, L. (2001). Random forests. *Mach. Learn.* **45**, 5–32. <https://doi.org/10.1023/A:1010933404324>.
  106. Choudhary, K., DeCost, B., and Tavazza, F. (2018). Machine learning with force-field-inspired descriptors for materials: Fast screening and mapping energy landscape. *Phys. Rev. Mater.* **2**, 083801. <https://doi.org/10.1103/PhysRevMaterials.2.083801>.
  107. Goodall, R.E.A., and Lee, A.A. (2020). Predicting materials properties without crystal structure: Deep representation learning from stoichiometry. *Nat. Commun.* **11**, 6280. <https://doi.org/10.1038/s41467-020-19964-7>.
  108. Chicco, D., and Jurman, G. (2020). The advantages of the Matthews correlation coefficient (MCC) over F1 score and accuracy in binary classification evaluation. *BMC Genom.* **21**, 6–13. <https://doi.org/10.1186/s12864-019-6413-7>.
  109. Bartel, C.J., Trewartha, A., Wang, Q., Dunn, A., Jain, A., and Ceder, G. (2020). A critical examination of compound stability predictions from machine-learned formation energies. *npj Comput. Mater.* **6**, 97. <https://doi.org/10.1038/s41524-020-00362-y>.
  110. Yamada, H., Liu, C., Wu, S., Koyama, Y., Ju, S., Shiomi, J., Morikawa, J., and Yoshida, R. (2019). Predicting materials properties with little data using shotgun transfer learning. *ACS Cent. Sci.* **5**, 1717–1730. <https://doi.org/10.1021/acscentsci.9b00804>.
  111. Jha, D., Choudhary, K., Tavazza, F., Liao, W.K., Choudhary, A., Campbell, C., and Agrawal, A. (2019). Enhancing materials property prediction by leveraging computational and experimental data using deep transfer learning. *Nat. Commun.* **10**, 5316. <https://doi.org/10.1038/s41467-019-13297-w>.
  112. Zhang, Y., and Ling, C. (2018). A strategy to apply machine learning to small datasets in materials science. *npj Comput. Mater.* **4**, 25. <https://doi.org/10.1038/s41524-018-0081-z>.
  113. Himanen, L., Jäger, M.O., Morooka, E.V., Federici Canova, F., Ranawat, Y.S., Gao, D.Z., Rinke, P., and Foster, A.S. (2020). DScribe: Library of descriptors for machine learning in materials science. *Comput. Phys. Commun.* **247**, 106949. <https://doi.org/10.1016/j.cpc.2019.106949>.
  114. Langer, M.F., Goeßmann, A., and Rupp, M. (2022). Representations of molecules and materials for interpolation of quantum-mechanical simulations via machine learning. *npj Comput. Mater.* **8**, 41. <https://doi.org/10.1038/s41524-022-00721-x>.
  115. Jia, R., and Liang, P. (2017). Adversarial Examples for Evaluating Reading Comprehension Systems. Preprint at arXiv. <https://doi.org/10.48550/arXiv.1707.07328>.
  116. DeYoung, J., Jain, S., Rajani, N.F., Lehman, E., Xiong, C., Socher, R., and Wallace, B.C. (2019). ERASER: A benchmark to evaluate rationalized NLP models. Preprint at arXiv. <https://doi.org/10.48550/arXiv.1911.03429>.
  117. Simonyan, K., Vedaldi, A., and Zisserman, A. (2013). Deep inside convolutional networks: Visualising image classification models and saliency maps. Preprint at arXiv. <https://doi.org/10.48550/arXiv.1312.6034>.
  118. Sundararajan, M., Taly, A., and Yan, Q. (2017). Axiomatic attribution for deep networks. In *International conference on machine learning (PMLR)*, pp. 3319–3328.

119. Ribeiro, M.T., Singh, S., and Guestrin, C. (2016). Why should i trust you?" Explaining the predictions of any classifier. In Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining, pp. 1135–1144. <https://doi.org/10.1145/2939672.2939778>.
120. Li, J., Monroe, W., and Jurafsky, D. (2016). Understanding neural networks through representation erasure. Preprint at arXiv. <https://doi.org/10.48550/arXiv.1612.08220>.
121. Krenn, M., Pollice, R., Guo, S.Y., Aldeghi, M., Cervera-Lierta, A., Friederich, P., dos Passos Gomes, G., Häse, F., Jinich, A., Nigam, A., et al. (2022). On scientific understanding with artificial intelligence. *Nat. Rev. Phys.* *4*, 761–769. <https://doi.org/10.1038/s42254-022-00518-3>.
122. van Dis, E.A.M., Bollen, J., Zuidema, W., van Rooij, R., and Bockting, C.L. (2023). ChatGPT: five priorities for research. *Nature* *614*, 224–226. <https://doi.org/10.1038/d41586-023-00288-7>.
123. Muggleton, S.H., Schmid, U., Zeller, C., Tamaddoni-Nezhad, A., and Besold, T. (2018). Ultra-strong machine learning: comprehensibility of programs learned with ILP. *Mach. Learn.* *107*, 1119–1140. <https://doi.org/10.1007/s10994-018-5707-3>.
124. Klimeš, J., Bowler, D.R., and Michaelides, A. (2009). Chemical accuracy for the van der Waals density functional. *J. Phys. Condens. Matter* *22*, 022201. <https://doi.org/10.1088/0953-8984/22/2/022201>.
125. Tran, F., and Blaha, P. (2009). Accurate band gaps of semiconductors and insulators with a semilocal exchange–correlation potential. *Phys. Rev. Lett.* *102*, 226401. <https://doi.org/10.1103/PhysRevLett.102.226401>.
126. Choudhary, K., Zhang, Q., Reid, A.C.E., Chowdhury, S., Van Nguyen, N., Trautt, Z., Newrock, M.W., Congo, F.Y., and Tavazza, F. (2018). Computational screening of high-performance optoelectronic materials using OptB88vdW and TB-mBJ formalisms. *Sci. Data* *5*, 180082. <https://doi.org/10.1038/sdata.2018.82>.
127. Choudhary, K., Bercx, M., Jiang, J., Pachter, R., Lamoen, D., and Tavazza, F. (2019). Accelerated discovery of efficient solar cell materials using quantum and machine-learning methods. *Chem. Mater.* *31*, 5900–5908. <https://doi.org/10.1021/acs.chemmater.9b02166>.
128. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. (2019). Pytorch: An imperative style, high-performance deep learning library. *Adv. Neural Inf. Process. Syst.* *32*.
129. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. (2011). Scikit-learn: Machine learning in Python. *J. Mach. Learn. Res.* *12*, 2825–2830.
130. Loshchilov, I., and Hutter, F. (2017). Decoupled weight decay regularization. Preprint at arXiv. <https://doi.org/10.48550/arXiv.1711.05101>.
131. Wang, M., Zheng, D., Ye, Z., Gan, Q., Li, M., Song, X., Zhou, J., Ma, C., Yu, L., Gai, Y., et al. (2019). Deep graph library: A graph-centric, highly-performant package for graph neural networks. Preprint at arXiv. <https://doi.org/10.48550/arXiv.1909.01315>.
132. Wu, Y., Schuster, M., Chen, Z., Le, Q.V., Norouzi, M., Macherey, W., Krikun, M., Cao, Y., Gao, Q., Macherey, K., et al. (2016). Google's neural machine translation system: Bridging the gap between human and machine translation. Preprint at arXiv. <https://doi.org/10.48550/arXiv.1609.08144>.
133. Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., et al. (2020). Transformers: State-of-the-art natural language processing. In Proceedings of the 2020 conference on empirical methods in natural language processing: system demonstrations, pp. 38–45. <https://doi.org/10.18653/v1/2020.emnlp-demos.6>.
134. Kokhlikyan, N., Miglani, V., Martin, M., Wang, E., Alsallakh, B., Reynolds, J., Melnikov, A., Kliushkina, N., Araya, C., Yan, S., et al. (2020). Captum: A unified and generic model interpretability library for pytorch. Preprint at arXiv. <https://doi.org/10.48550/arXiv.2009.07896>.
135. Shapley, L.S. (1953). *A Value for N-Person Games* (Princeton University Press).
136. Owen, G. (1977). Values of games with a priori unions. In *Mathematical economics and game theory: Essays in honor of Oskar Morgenstern* (Springer), pp. 76–88. [https://doi.org/10.1007/978-3-642-45494-3\\_7](https://doi.org/10.1007/978-3-642-45494-3_7).