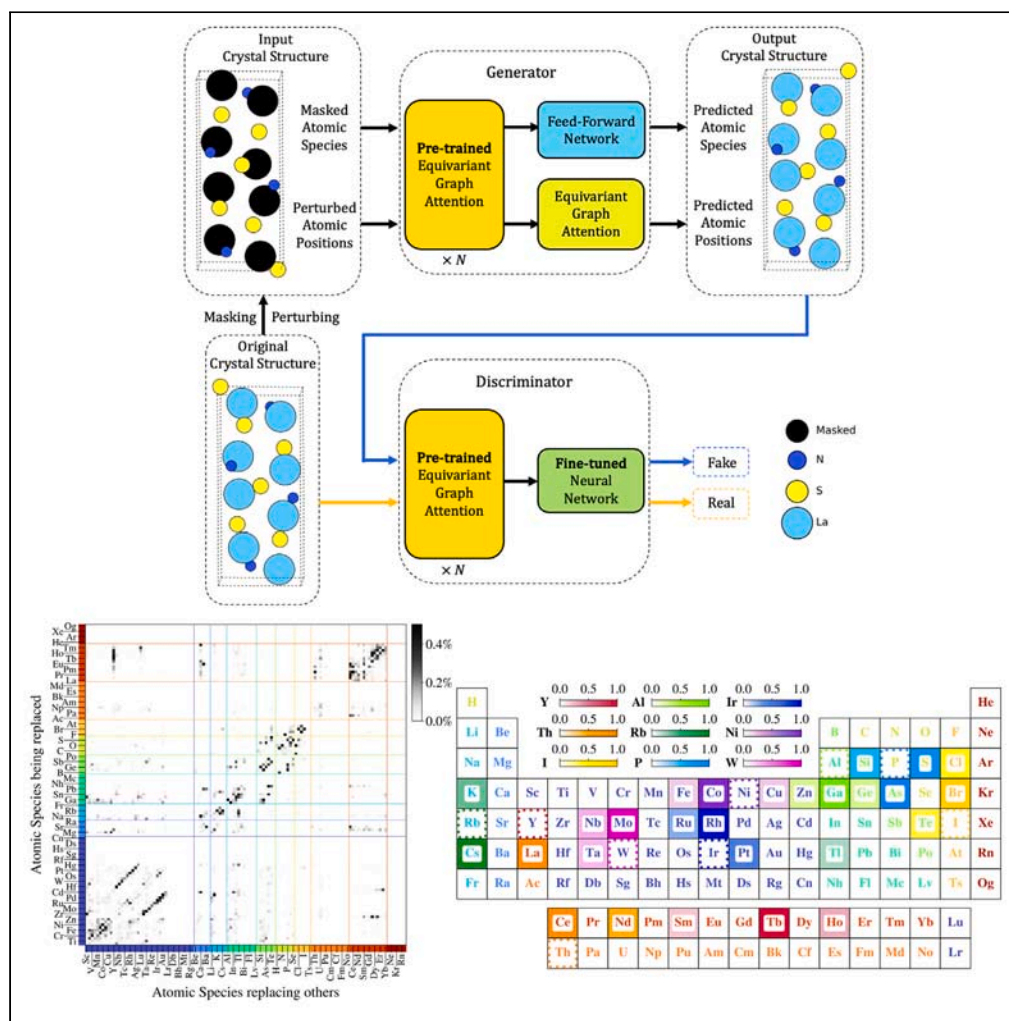


Article

Self-supervised generative models for crystal structures



Fangze Liu,
Zhantao Chen,
Tianyi Liu, Ruyi
Song, Yu Lin,
Joshua J. Turner,
Chunjing Jia

fangzel@stanford.edu (F.L.)
zhantao@stanford.edu (Z.C.)
joshuat@slac.stanford.edu
(J.J.T.)
chunjing@phys.ufl.edu (C.J.)

Highlights

Established a unified
framework for material
generation and property
prediction

Trained the network with
self-supervised and
adversarial learning for
better reliability

The trained model
demonstrated the
extraction of intrinsic
material information

The trained model
revealed correlations
between materials through
generative modeling

Article

Self-supervised generative models for crystal structures

Fangze Liu,^{1,2,6,*} Zhantao Chen,^{2,3,6,7,*} Tianyi Liu,^{2,4} Ruyi Song,² Yu Lin,² Joshua J. Turner,^{2,3,*} and Chunjing Jia^{2,5,*}

SUMMARY

Inspired by advancements in natural language processing, we utilize self-supervised learning and an equivariant graph neural network to develop a unified platform for training generative models capable of generating inorganic crystal structures, as well as efficiently adapting to downstream tasks in material property prediction. To mitigate the challenge of evaluating the reliability of generated structures during training, we employ a generative adversarial network (GAN) with its discriminator being a cost-effective reliability evaluator, significantly enhancing model performance. We demonstrate the utility of our model in optimizing crystal structures under predefined conditions. Without external properties acquired experimentally or numerically, our model further displays its capability to help understand inorganic crystal formation by grouping chemically similar elements. This paper extends an invitation to further explore the scientific understanding of material structures through generative models, offering a fresh perspective on the scope and efficacy of machine learning in material science.

INTRODUCTION

Structure and property predictions of crystalline materials have been a long-standing and central focus in condensed matter physics and material sciences. Recent advancements have demonstrated the efficacy of machine learning techniques in predicting various materials properties, including electronic topology, thermodynamic properties, mechanical moduli, etc.^{1–9} Furthermore, there has been a growing focus toward applying generative models inspired from computer vision, such as generative adversarial networks (GANs),^{10–12} diffusion models,^{13–15} and variational autoencoders (VAEs)^{14,16} for crystal structure generation. GANs, which involve adversarial training of two neural networks, have been applied to a limited range of materials, like specific compositional families¹⁰ and two-dimensional materials.^{11,12} VAEs, which encode material representation into a compressed latent space and generate materials by decoding sampled latent codes in the space, is a favored architecture that can be effectively combined with other training methods like GAN¹⁰ and diffusion processes,¹⁴ though they demand additional efforts for effective latent space interpretation and utilization. Diffusion models, known as a physics-induced models and celebrated in computer vision,¹⁷ have seen successful adaptation for crystal generation,¹⁴ a trend underscored by a recent surge of studies.^{15,18} Instead of using generative models, the combination of simple element substitutions and density functional theory (DFT)^{19,20} hints at the benefits of incorporating DFT into the generative learning frameworks for enhancing crystal generation.

One fundamental question is focused on how to understand the connection between atomic structures and the properties of materials. Machine learning techniques have proven powerful in predicting the latter, given the former, especially as the dataset of stable crystal structures is currently much larger than the dataset of stable structures with properties obtained either numerically or experimentally. However, predicting atomic structures presents considerably greater difficulties, including the vast structure design space that exists for materials discovery and the lack of suitable evaluation metrics for generated materials. This situation is similar to the challenges faced in natural language processing (NLP), as the volume of unlabeled textual data is much larger than that of labeled question-answer data. Another level of similarity is observed in the distribution of atomic species in crystal structures compared to the distribution of vocabulary in natural language data²¹ (see Figure 1A). As studies have found that the emergent learning ability of large transformer-based language models²² is attributed to the skewed and long-tailed distributional properties of the training data,²³ this suggests that the established strategies adopted in training large language models (LLMs) could potentially be successful in training machine learning models for materials science. Comparing to the costly in-context learning,²⁴ self-supervised learning is a more approachable training method^{25–27}: a neural network is pre-trained by large volumes

¹Department of Physics, Stanford University, Stanford, CA 94305, USA

²Stanford Institute for Materials and Energy Sciences, SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA

³Linac Coherent Light Source, SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA

⁴Department of Chemistry, Stanford University, Stanford, CA 94305, USA

⁵Department of Physics, University of Florida, Gainesville, FL 32611, USA

⁶These authors contributed equally

⁷Lead contact

*Correspondence: fangzel@stanford.edu (F.L.), zhantao@stanford.edu (Z.C.), joshuat@slac.stanford.edu (J.J.T.), chunjing@phys.ufl.edu (C.J.)
<https://doi.org/10.1016/j.isci.2024.110672>



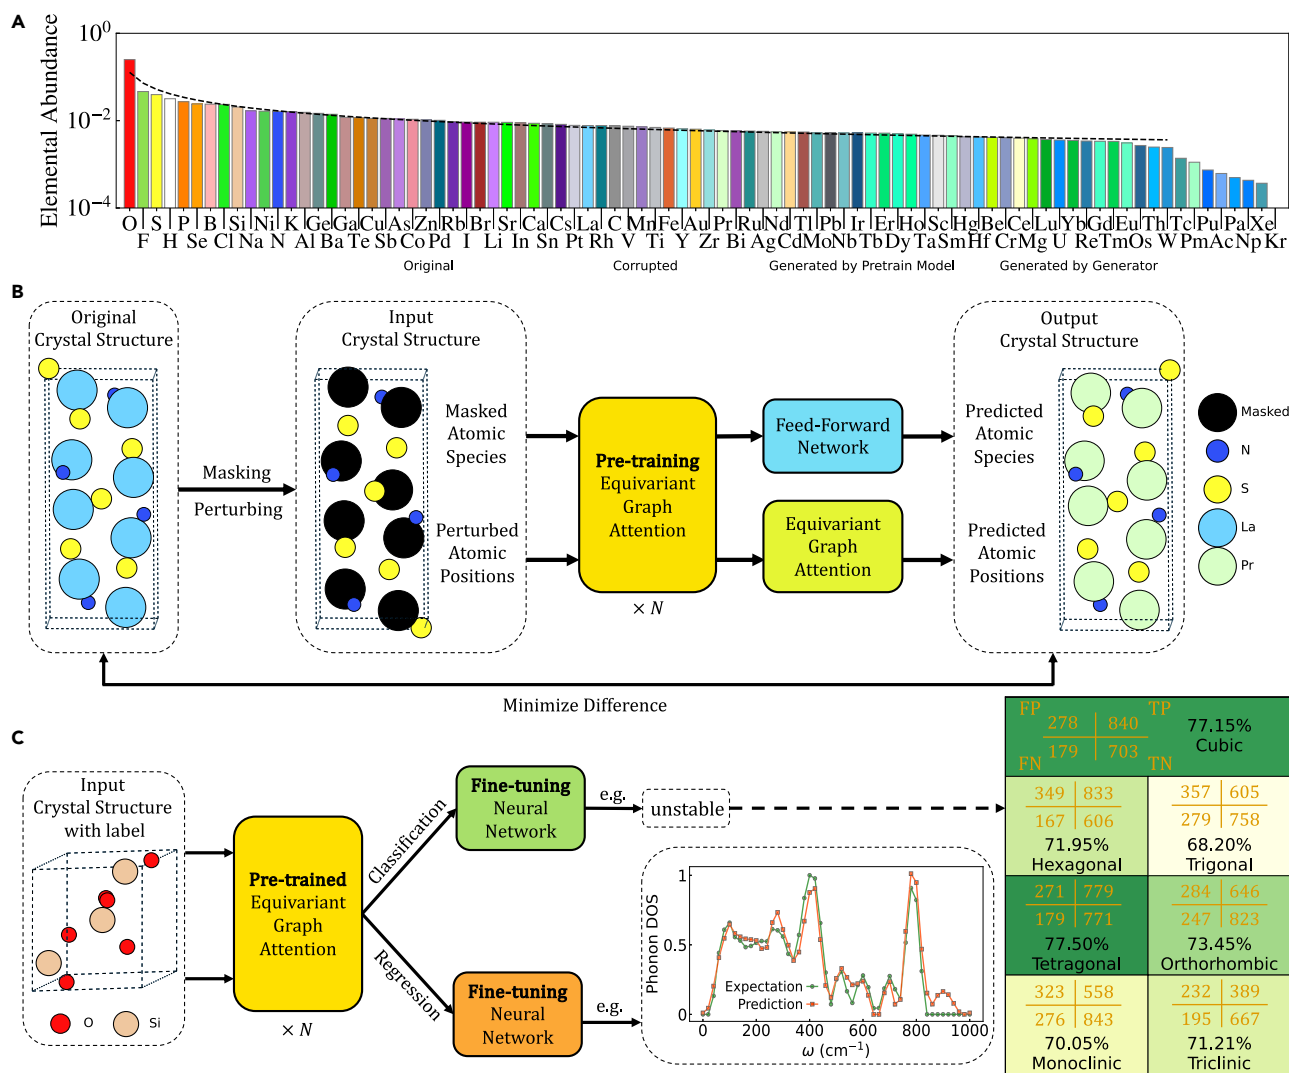


Figure 1. Self-supervised learning in crystal structure generation

(A) Training samples show a power law elemental abundance distribution, $p(X = x) \propto 1/x^\alpha$, where X is the rank of an element and $\alpha = 0.8163$ with excluding the the eight least frequently occurring elements. This resulting distribution (dashed line) resembles the natural language token distribution, typically with an exponent close to 1,²³ supporting the utilization of self-supervised learning in crystal generation.

(B) The self-supervised pre-training employs an equivariant graph attention transformer architecture, EquiformerV2,²⁹ processing contaminated structures produced by masking parts of the atoms and perturbing atomic positions in the original stable crystal structures, and being trained on tasks to predict the masked atoms and restore atomic positions. Primitive lattice vectors are predefined to compute the edge distance embedding, serving as input for Equiformer.

(C) The pre-trained model is well equipped for various supervised down-stream tasks through the fine-tuning procedure. During fine-tuning, the pre-trained model is concatenated with a randomly initialized feedforward neural network and all model parameters undergo fine-tuning. To showcase the model's classification ability, an example of predicting crystal stability, using a labeled training set based solely on cubic system, is illustrated. The accompanying diagram displays classification results (true positives, true negatives, false negatives, and false positives in a clockwise order), with accuracy indicated by percentage numbers and reflected by colors. An example of the regression tasks, predicting phonon density of states (DOS), is also depicted. The training set is sourced from the Materials Project. More details are provided in [STAR methods](#).

of unlabeled and augmented data, like randomly masking tokens or shuffling the order of sentences.²⁸ The pre-trained model can capture the nuanced patterns and structures in training data, and further be fine-tuned on task-specific labeled data and reduce the reliance on expensive labeled datasets.

Taking inspiration from NLP, we present a machine learning framework for material generation and properties prediction, empowered by an efficient pre-training strategy without the need of human knowledge as a prerequisite for labeling. Analogous to a typical training procedure for NLP models, where fill-in-the-blank and sentence-arranging exercises are utilized, we design a self-supervised training procedure for models based on the state-of-the-art transformer-based equivariant graph neural network, EquiformerV2.²⁹ During pre-training, the model

takes contaminated crystal structures as inputs, in which atomic species and positions are randomly masked and perturbed, and learns to reconstruct the complete and noiseless structures. In tests with crystal structures previously unseen by the model, the pre-trained model has shown a preliminary ability to identify local optimal solutions from incomplete structures. Hence, when provided with a basic design for the structure of novel materials, our model can produce a completion plan, generating the most likely stable crystal structures under given conditions. In addition, we also demonstrate the adaptability of our pre-trained model for downstream regression and classification tasks through supervised fine-tuning.

The challenge, however, lies in the lack of a critic to evaluate the generated structures efficiently. We consider using the actor-critic learning framework to guide generative training, similar to the application of reinforcement learning from human feedback (RLHF) in NLP.³⁰ While numerical methods like DFT calculations are potential options for the critic's role, leveraging GANs offers a significantly more computationally efficient solution for early-stage training. More importantly, our goal extends beyond mere generative tasks. We aim for models to reveal the intrinsic information embedded within material structures. For example, our model serves as a conditional probabilistic model for investigating the likelihood of various compositions being stable under given crystallographic conditions, so it offers a probabilistic insight into the nature of crystals and delivers richer information compared to other unconditional probabilistic models derived through data mining.^{31,32} Consequently, we incorporate the GAN architecture into our training pipeline to fully exploit existing data without relying on external information, such as DFT-calculated stability labels, and thereby circumvent the limitations imposed by human preconceptions.

In this paper, we explore the application of self-supervised learning to crystal structure generation and demonstrate that the incorporation of a discriminative model can enhance the reliability of the generated crystal structures with minimal additional effort. This enhancement is evidenced by a comprehensive comparison between the outcomes of self-supervised learning and its combination with GAN. Our methodology is unique as it relies solely on unlabeled crystal structure data, offering a data- and computation-efficient training strategy for early-stage training. Furthermore, this approach paves the way for a first-principles understanding of the intrinsic information hidden in existing material structures, thereby serving as an invitation for further exploration into the analysis of material structures through generative models.

RESULTS

Pre-training for crystal generation

To obtain a generative model, we integrate the concepts of masked training and machine learning denoising into equivariant graph neural networks. Given a stable crystal structure as a valid training sample, we prepare an imperfect input structure by masking a portion of the atoms, i.e., setting their atomic numbers to zero, and perturbing the equilibrium atomic positions. In particular, the incomplete input structures are generated through the following operations: (1) masking all atoms of a randomly selected species or randomly masking 15% of all atoms in the unit cell, regardless of species, where the choice between these masking strategies is made randomly; and (2) adding random displacements to the positions of all atoms, including the masked ones. The primitive lattice vectors and periodicity of the input structures are predefined to compute the edge distance embedding, serving as input for the generative model. As illustrated in Figure 1B, the model is designed to reconstruct the complete structures from the given inputs by predicting the atomic species and positions in separate task-specific layers, and the discrepancies between the reconstructed structures and the pristine structures are used to update model weights.

Specifically, the model incorporates the main structure of the EquiformerV2,²⁹ an equivariant graph neural network with the attention mechanism, and two auxiliary, shallow neural networks that further map the equivariant features to the desired atomic species and position information. The equivariant backbone model efficiently represents crystal structures with all symmetries preserved and provides informative embedding for the subsequent layers. To train the model, we adopt a hybrid loss function that combines the negative log likelihood for atomic species predictions and the mean squared error (MSE) for position predictions. Further details about the model and its training process are provided in STAR methods.

Fine-tuning for downstream tasks

The pre-trained backbone model can serve as a versatile pre-trained network block adaptable for various downstream property-prediction tasks by connecting with additional shallow layers. For each specific supervised task, this pre-trained model is integrated with a feedforward layer, which is randomly initialized and is adapted to convert the output of the pre-trained model to meet requirements of the new task. During fine-tuning, both the pre-trained model and the feedforward layer are trained on a task-specific dataset. This approach allows the integrated model to be rapidly and efficiently tailored for a variety of tasks, offering significant advantages over training a new model from scratch.

To illustrate the capability of performing classification as an exemplary downstream task, we showcase the model performance on predicting the stability of a material structure in the inset table of Figure 1C. Here, we use the stability labels from the Materials Project database determined by the convex hull. Notably, the fine-tuning network is trained on structures with cubic lattices only and reaches a commendable accuracy of 77.15%. It further generalizes reasonably well to diverse lattice types, yielding close accuracy for tetragonal lattices and around 70% accuracies for other lattice types that the model has not been trained with. Such fine-tuning ideas can also be extended to regression tasks, such as predicting the phonon density of states (DOS), as shown in the lower part of Figure 1C. A diverse set of other regression tasks including predictions of Fermi energy, bulk moduli, and shear moduli are further discussed in STAR methods. It is worth mentioning that, for all the demonstrated downstream tasks, the material structures—specifically, the atomic species and positions—are the only inputs. We intentionally exclude any additional atomic properties, such as covalent radius and electronegativity, to explore the feasibility of predicting physical properties based solely on first principles. The remarkable prediction performances on most examined tasks indicate the promising potential of our approach.

GAN for more reliable structure generation

Through self-supervised learning on a dataset consisting of over 20,000 diverse, stable crystal structures from the Materials Project,³³ the pre-trained model displays proficiency in accurately predicting atom types and positions in the majority of cases, with a few examples displayed in Figure 3 and additional examples available in STAR methods. However, the model with the architecture presented in Figure 1B encounters limitations. For example, the model faces challenges in accurately reconstructing species that occur infrequently within our dataset. It tends to replace the masked species with alternatives possessing similar chemical properties, as depicted in Figure 3B, thereby leading to a lack of compositional validity in some outputs.

To mitigate the aforementioned challenges in generating physically valid structures, it becomes essential to incorporate a mechanism for evaluating the generated structures. Utilizing human pre-knowledge, such as DFT, to train discriminator models is a common practice and has led to successful outcomes in some recent studies.^{19,20} However, it can become exceedingly demanding in terms of computational resources, as DFT calculations for each crystal structure can take several minutes to even hours, depending on the number and species of atoms and their types in the unit cell. Considering the need to evaluate as many generated crystals as training samples in each training epoch, a DFT-based evaluation mechanism is thus impractical or at least inefficient. Moreover, to assess the ability of predicting crystal structures and physical properties based solely on material structures—free from the bias of human preconceptions—we also intentionally exclude any additional atomic properties observed by experiments or numerical methods. It highlights the pressing need for developing more computationally efficient and unbiased evaluation methods.

Given these constraints, GANs emerge as a more practical solution. Following the GAN framework illustrated in Figure 2, we combine our pre-trained generative model (Figure 1B) with a discriminative model that has been initialized using fine-tuned parameters for stability predictions (Figure 1C), motivated by the similarities between stability classification and the discriminator's role. Specifically, the generator is tasked with reconstructing material structures from incomplete and perturbed input structures. The output structures are further passed into the discriminative model along with the pristine structures. The goal of the generator is to produce seemingly realistic materials that can pass the test by the discriminator, while the discriminator tries to tell the difference of generated (fake) structures from the real ones.

To provide an intuitive understanding of our models' performance, we provide representative examples of the original material structures, contaminated inputs, and the corresponding outputs generated by different models in Figure 3. The generative model, trained under the GAN framework, demonstrates visibly improved performance. It is particularly evident when atoms of a less common species are masked: the GAN generator outperforms the pre-trained model in reasonably repositioning atoms and providing more reliable predictions for the masked species. While evaluating the enhancement offered by GANs over pre-trained model purely from a visual perspective is not persuasive, we quantitatively assess their differences based on three metrics: validity, similarity, and novelty.

Validity

We conduct an analysis of the validity of generated crystal structures using two fundamental criteria commonly employed in the field^{14–16}: structural validity and compositional validity. *Structural validity* assesses whether the minimum Euclidean distance between atoms is greater than a threshold of 0.5 Å, ensuring appropriate atomic spacing. *Compositional validity* is determined by verifying if the chemical composition maintains a neutral charge, which is calculated by the Semiconducting Materials by Analogy and Chemical Theory.³⁴

As shown in Figure 4A, when the noise level in the testing inputs matches that of the training inputs, the GAN generator, though not outperforming the pre-trained model in constructing structurally valid crystals, demonstrates enhanced robustness in generating compositionally valid structures. Despite all generative models being trained on inputs with only 15% of atoms masked, the GAN approach displays remarkable extrapolation capability, especially in scenarios where 30% atoms are randomly masked (mask type γ).

Interestingly, the GAN generator trained with a higher noise level ($\mathcal{G}_{0.2}$) performs similarly to other models on the test set with a lower noise level ($\tilde{\sigma}_{\text{noise}} = 0.1$); however, when test samples include higher noise levels, GAN generators outperform the pre-trained model in creating compositionally and structurally valid crystals, specially the structural validity score of $\mathcal{G}_{0.2}$ is distinctly higher than that of \mathcal{G}_{pre} . $\mathcal{G}_{0.2}$ uses a training set with $\tilde{\sigma}_{\text{noise}} = 0.2$ during GAN training, but since it utilizes the model parameters of the pre-trained model \mathcal{G}_{pre} , it has been effectively exposed to training set with different noise levels. This observation suggests that diversifying the training set by including variations in $\tilde{\sigma}_{\text{noise}}$ may enhance the model's versatility.

Similarity

We introduce two metrics designed to evaluate the effectiveness of generative models in accurately reconstructing the crystal structures of ground truth materials, namely compositional similarity and structural similarity. *Compositional similarity* quantifies the probability of generative models in reproducing the same compositions as the originals (see Equation 5). *Structural similarity* reflects their accuracy in replicating the atomic positions in alignment with their original counterparts, scaled by the noise introduced to the inputs (as defined by Equation 4).

Figure 4B shows that the distribution of compositional and structural similarities approximates a Gaussian distribution, with the mean values locating at the standard deviation of the noise distribution introduced to the training samples, $\sigma_{\text{noise}} = \tilde{\sigma}_{\text{noise}} \times \min(\text{edges})$. As the noise level of the test set increases, the models quickly become incapable of reverting the compositions and structures to their originals. This inability could be attributed to the generated structures converging to other local optima, distinct from the original structures.

Figures 4D and 4E visualize the impact of the crystal size and the species to be unmasked on models' performance. It demonstrates that, despite the non-uniform distribution of the size of crystals (illustrated in STAR methods) and species abundance within the datasets (as shown in Figure 1A), the structural similarity scores for all models remain largely unaffected by variations in the number of atoms and the types of

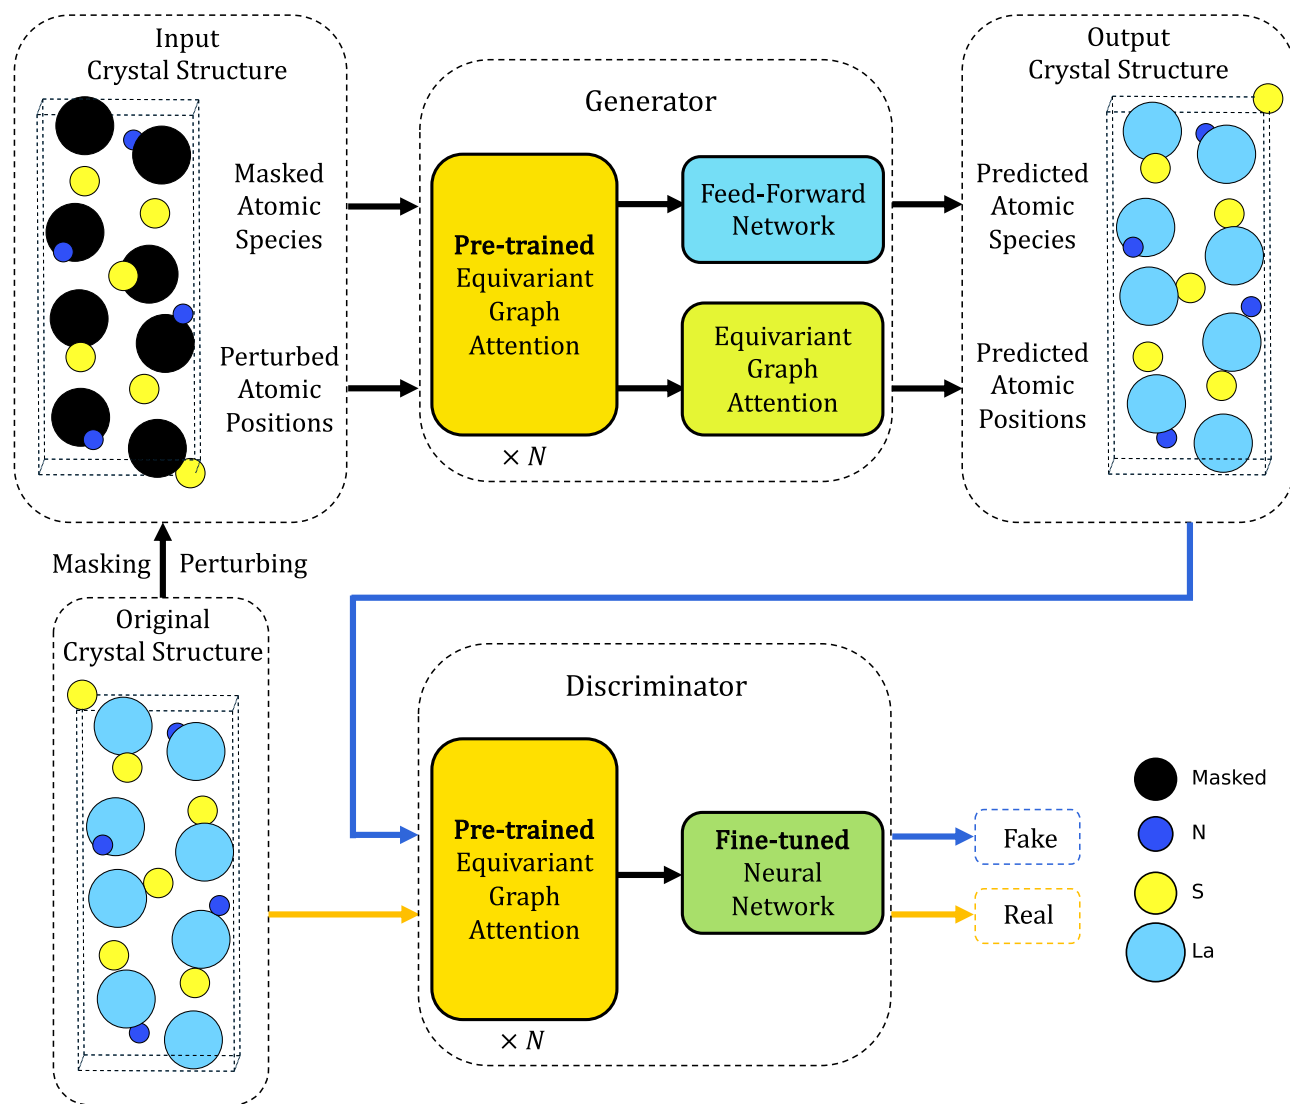


Figure 2. Generative adversarial network for crystal structure generation

We further adopt a GAN framework as a strategic enhancement to the pre-trained generative model. The discriminator, initialized with the fine-tuned stability-predicting model, is to distinguish between original stable and generated crystal structures, while the generator, initialized with the pre-trained model, aims to produce crystal structures indistinguishable from the real ones to the discriminator. All model parameters undergo fine-tuning during the procedure illustrated in this figure.

species masked. This observation suggests that the generative models exhibit significant robustness in reconstructing crystal structures, regardless of the difference in size and composition encountered during training.

DFT-based stability check

Another aspect of the performance of our model is the stability of the generated structures. Due to the relatively high computational cost and complexity of DFT simulations (especially for complex systems, e.g., large unit cells containing hundreds of atoms^{35–40} and systems containing polarized electrons^{41–43}), we randomly select a small subset of systems studied in this work (with various unit cell sizes) to conceptually show the relative stability of structures generated by different models.

Here we use the DFT total energy to feature the relative stability. A lower total energy value means a higher structural stability for systems with the same chemical composition. To eliminate the influence of unit cell size, we further divide the total energy values by the number of atoms in the unit cell. Figure 4C shows that structures produced by GAN generators usually exhibit a non-negligible lower total energy (per atom) ΔE , highlighting the effectiveness of GANs in constructing more stable structures in comparison to those generated by the pre-trained model. Additionally, the generated structure with the lowest total energy for each composition achieves the highest structural similarity score

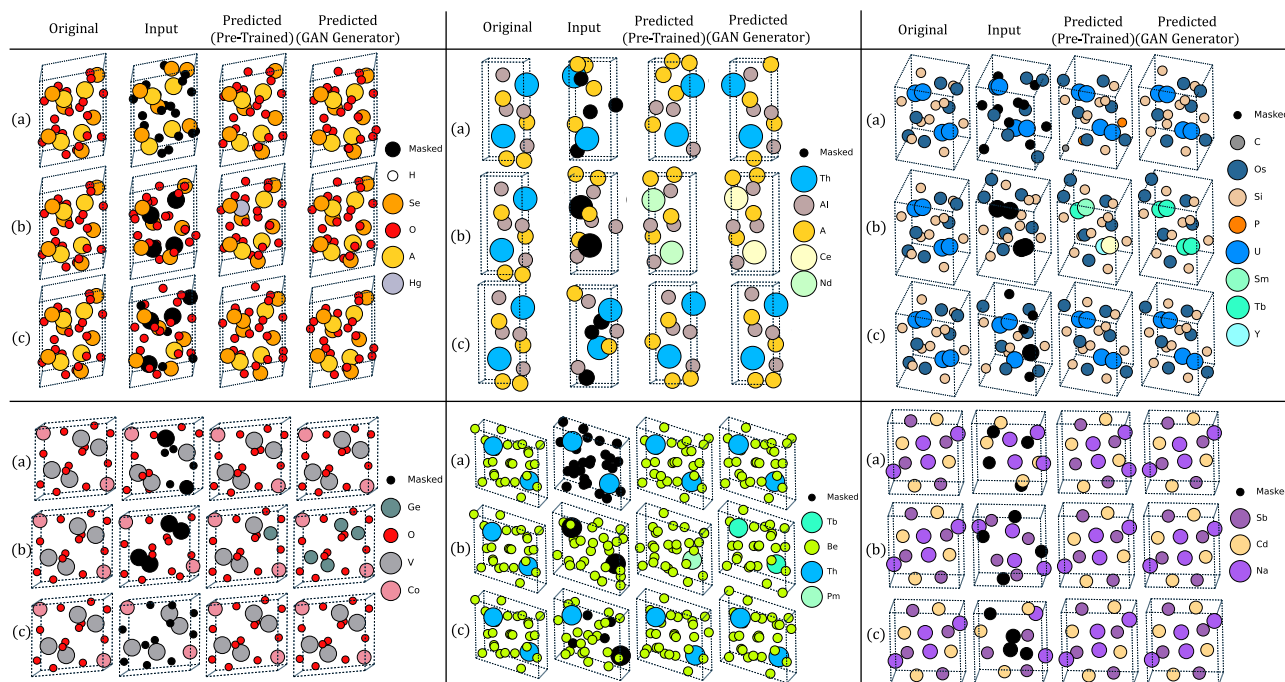


Figure 3. Generative/reconstructive capability illustration

(A) Showcase the capability of the generators in finding the optimal crystal structures from inputs with *frequent occurring species* being masked. (B) Showcase the capability of the generators in finding the optimal crystal structures from inputs with *infrequent occurring species* being masked. (C) Showcase the capability of the generators in finding the optimal crystal structures from inputs with 30% *random atoms* being masked. Note that for all results shown in this figure, all positions are perturbed by noise following a normal distribution with a standard deviation of $\bar{\sigma}_{\text{noise}} \times \text{min}(\text{edges})$ ($\bar{\sigma}_{\text{noise}} = 0.2$). The pre-trained model is trained with $\bar{\sigma}_{\text{noise}} = 0.1$, and the GAN generator with $\bar{\sigma}_{\text{noise}} = 0.2$. More examples are provided in [STAR methods](#).

(refer to [STAR methods](#)). This observation suggests that structural similarity may act as a proxy for the structural stability suggested by DFT calculations.

Novelty

To gain deeper insights into the information encoded by generative models, as shown in [Figure 5](#), we plot the correlation between the species being replaced and those replacing them in the generated crystal structures. In this analysis, we only consider structures that differ in composition from their original counterparts and are compositionally and structurally valid. Our observations suggest that the model may automatically capture fundamental chemical properties from the self-supervised learning process without the needs of explicit labels, such as electron configurations, ionic radii, oxidation states, and the position of the element in the periodic table. For example, elements within the 4d series are preferred to be interchangeable with those in the 5d series within the same group, due to their similar electronic structures. Yttrium, for instance, is observed to be replaceable by terbium and holmium, as they share the same number of outer electrons, comparable atomic radii and oxidation states. Moreover, a common trend is noted where many elements exhibit a preference for replacement with others in close proximity on the periodic table (see [Figure 5B](#)). Such high level information provided by our pre-trained model can also be further utilized to help find novel stable compounds by narrowing down most promising candidate elements.

We further compare our method with simple elemental substitution method,³¹ which computes the likelihood of substituting species *A* with species *B*, denoted as $P(A \rightarrow B)$, through data mining. This method is particularly useful when the probability of one species being replaced by another remains constant, irrespective of different atomic and positional information within various crystals—although such scenario rarely applies. On the other hand, our generative model is essentially a more sophisticated probabilistic model, since it outputs the probability of specific atomic species at each masked site. Therefore, our model can offer conditional probabilities $P(A \rightarrow B | \text{condition})$ under diverse conditions, such as varying structures and compositions. We also find that the averaged probabilities over conditions given by our generative models bear a resemblance to those produced by simple substitution (see [STAR methods](#)).

DISCUSSION

In this study, we propose a machine learning framework that employs equivariant graph neural networks and utilizes the training strategies of reconstructing corrupted inputs ([Figure 1B](#)) and adversarial learning ([Figure 2](#)) to study the mechanisms underlying crystal structures

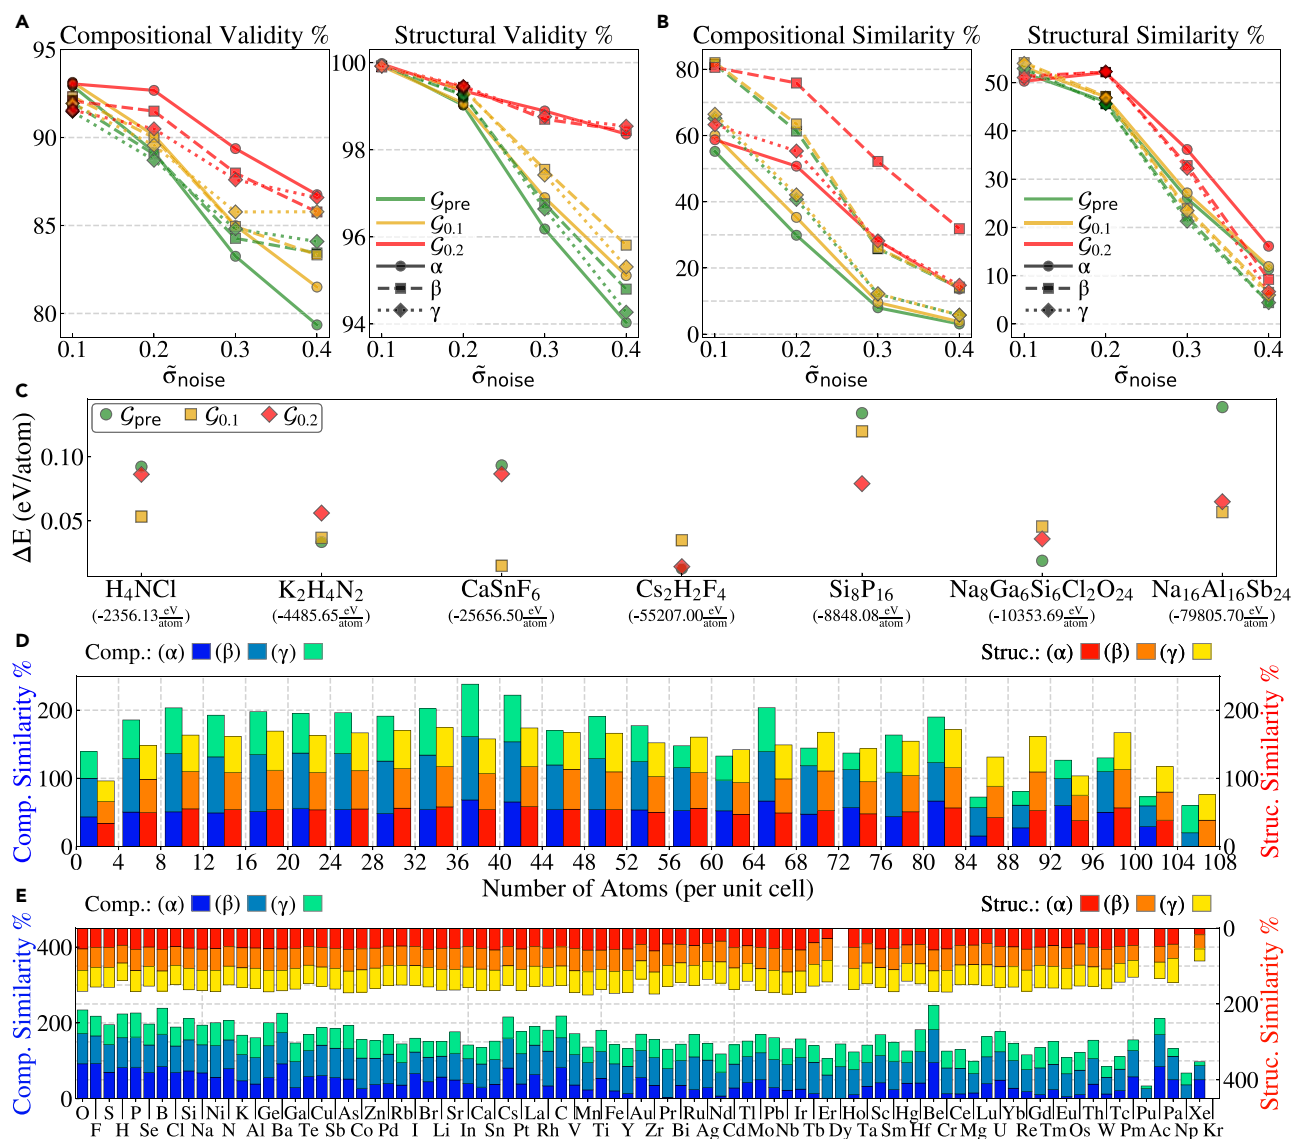


Figure 4. Performance evaluation

(A) The performance of generative models evaluated by *validity*. Model notation: G_{pre} represents the pre-trained generative model, trained with samples perturbed at $\bar{\sigma}_{noise} = 0.1$; $G_{0.1}$ and $G_{0.2}$ denote GAN generators trained with data at noise levels of $\bar{\sigma}_{noise} = 0.1$ and $\bar{\sigma}_{noise} = 0.2$, respectively.

(B) The performance of generative models evaluated by *similarity*.

(C) The performance of generative models evaluated by *DFT calculations*. Crystal structures, randomly selected from the test set and perturbed by positional noise with $\bar{\sigma}_{noise}$, are analyzed for energy difference (ΔE) compared to originals. The total energy per atom of originals is shown in brackets.

(D and E) Evaluate the impacts of the number of atoms per unit cell and the atomic species being masked on compositional similarity (blue-series histograms, left axis) and structural similarity (red-series histograms, right axis), using model $G_{0.2}$ and testing on samples with $\bar{\sigma}_{noise} = 0.2$. Note that the right axis in (E) is inverted. More detailed analyses are provided in [STAR methods](#). Positional noise for the test samples follows a Gaussian distribution with a standard deviation $\sigma_{noise} = \bar{\sigma}_{noise} \times \min(\text{edges})$. Masking strategies include the following: (α) masking all atoms of a specific species within each crystal structure, repeated for each species; (β) randomly selecting and masking 15% of atoms, repeated five times per structure; and (γ) a similar approach with 30% of atoms. Evaluation is performed on a test set augmented from 2,000 samples to approximately 6,500 for mask type (α) and 10,000 for mask type (β) and (γ).

formation. Our approach is designed to pursue three primary applications: (1) generating or reconstructing crystalline material structures, (2) predicting material properties, and (3) conducting behavioral and structural evaluations for the generative models trained exclusively on crystal structures.

Our model demonstrates a high capability in generating or reconstructing crystal structures: in case where, given any partial information about the original structure, the original remains the most optimal (stable) among all possibilities, our model achieves up to 80% accuracy in reconstructing the optimal composition and possesses a 60% efficiency in denoising the corrupted input toward the optimal structure (as

composing poetry, our approach focuses on training the model to learn crystal structures rather than directly instructing it to generate band structures.

In model evaluation, we assess the performance of the generative models on various test sets by examining aspects of validity, similarity, and novelty. We provide a detailed analysis of the model's capability across different input samples, focusing particularly on the influence of the number of atoms and the types of atomic species involved, with more features (such as symmetry groups) awaiting further exploration. One noteworthy observation is that the structural similarity, as defined in Equation 5, remains consistent across structures of varying sizes and compositions (see Figure 4). Additionally, we endeavor to examine whether crystal structures inherently contain all the information necessary for a comprehensive understanding of their formation, so our approach intentionally minimizes the reliance on pre-established human knowledge and assumptions. In attempting to analyze the generative model's underlying logic in selecting replacement elements, we find that the model may be capable of finding chemically similar species and providing novel data-driven insights related to the periodic table of elements (refer to Figure 5). Beyond behavioral evaluations, conducting structural evaluations, such as feature attribution, is a promising avenue for future research.

Our work opens the door for significant future advancements. Drawing inspiration from NLP, we implement an efficient self-supervised training method, offering a comprehensive framework for crystal generation and property prediction. To overcome the lack of an accessible supervisor, we demonstrate that incorporating a GAN framework enables the generator to produce more reliable outcomes. This approach, free from prior knowledge, further permits the investigation of information intrinsically encoded within crystal structures. Therefore, our study serves as an invitation for deeper exploration into understanding materials from first principles using generative machine learning techniques.

Limitations of the study

Although our work marks a step forward, it encounters limitations. (1) Achieving "large crystal models" capable of accurately predicting crystal structures requires a more comprehensive dataset and larger model parameters. (2) Currently, our models have not been trained on dataset that include perturbed lattice vectors, but this limitation can be readily addressed by including a feedforward network for lattice vector prediction into the generator. (3) The dataset used for training exclusively contains 3D crystals, which limits our model's ability to predict crystals of lower dimensionality. (4) Our generative model is presently confined to cases where the basic design of the desired crystal structure should be predefined. In other words, unlike other diffusion-based studies that aim for optimal structures with given compositions, our current work practically focuses on finding the optimal composition based on a basic design of the structure. However, our training framework is adaptable for structural optimization by introducing higher noise levels. Employing a diffusion process as the generator within the GAN framework improves the outcomes in the absence of specific structural designs. (5) Training models on more challenging tasks may lead to better outcomes. For example, we could delete atoms and their positions, or add extra ones, enabling the model to decide whether to remove or generate atoms; we could also randomly swap the positions of some atoms, challenging the model to correct these alterations. Such operations can potentially improve the performance of generative models. However, due to limited computational resource, our investigations focus on masked atoms and perturbed atomic positions. (6) Analogous to enhancements of language models through RLHF, for a thorough examination of the generated structures and a further refinement for practical use, incorporating human insights becomes an inevitable next step. (7) To achieve the generation of materials with constrained properties, our model needs to be further fine-tuned using property-specific datasets, though this is primarily a consideration under resource limitations. We expect the ultimate objective as conditioning the generation process by text representations and facilitating training through zero-shot learning.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- RESOURCE AVAILABILITY
 - Lead contact
 - Materials availability
 - Data and code availability
- METHOD SUMMARY
 - Model setup
 - Loss functions
 - Related metrics
 - Motivations and discussions
 - Self-supervised pre-training stage
 - Supervised fine-tuning stage
 - Generative adversarial network stage
 - Comparative analysis of generative capability

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.isci.2024.110672>.

ACKNOWLEDGMENTS

We thank the insightful discussions with Cheng Peng, Wei Shao, Jay Qu, and Minkyung Han. This work is supported by the U.S. Department of Energy, Office of Science, Basic Energy Sciences under award No. DE-SC0022216. Y.L. acknowledges the support by the U.S. Department of Energy, Laboratory Directed Research and Development program at SLAC National Accelerator Laboratory, under contract No. DE-AC02-76SF00515. This research used computational resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231.

AUTHOR CONTRIBUTIONS

F.L. and Z.C. contributed equally to this research. F.L. conceived and designed the research with input from Z.C. Z.C. built the pre-train model and collected dataset. F.L. refined, expanded, and trained all models, performed experiments, and analyzed and visualized the results, with support from Z.C. and T.L. R.S. performed DFT calculations. Z.C., F.L., and J.J.T. wrote the manuscript with feedback and contributions from all authors. C.J., J.J.T., and Y.L. supervised this research.

DECLARATION OF INTERESTS

The authors declare no competing interests.

DECLARATION OF GENERATIVE AI AND AI-ASSISTED TECHNOLOGIES IN THE WRITING PROCESS

During the preparation of this work the authors used the large language model ChatGPT by OpenAI to refine the language and enhance the readability of this paper. All content was independently written by the authors before utilizing this tool, and the authors thoroughly reviewed and edited the content after its use, taking full responsibility for the content of the publication.

Received: January 12, 2024

Revised: March 25, 2024

Accepted: August 1, 2024

Published: August 6, 2024

REFERENCES

- Andrejevic, N., Andrejevic, J., Bernevig, B.A., Regnault, N., Han, F., Fabbri, G., Nguyen, T., Drucker, N.C., Rycroft, C.H., and Li, M. (2022). Machine-Learning Spectral Indicators of Topology. *Adv. Mater.* 34, 2204113. <https://doi.org/10.1002/adma.202204113>.
- 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>.
- Choudhary, K., DeCost, B., Chen, C., Jain, A., Tavazza, F., Cohn, R., Park, C.W., Choudhary, A., Agrawal, A., Billinge, S.J.L., et al. (2022). Recent advances and applications of deep learning methods in materials science. *NPJ Comput. Mater.* 8, 59. <https://doi.org/10.1038/s41524-022-00734-6>.
- Kong, S., Ricci, F., Guevarra, D., Neaton, J.B., Gomes, C.P., and Gregoire, J.M. (2022). Density of states prediction for materials discovery via contrastive learning from probabilistic embeddings. *Nat. Commun.* 13, 949. <https://doi.org/10.1038/s41467-022-28543-x>.
- Liu, Y., Tan, X., Liang, J., Han, H., Xiang, P., and Yan, W. (2023). Machine learning for perovskite solar cells and component materials: key technologies and prospects. *Adv. Funct. Mater.* 33, 2214271. <https://doi.org/10.1002/adfm.202214271>.
- Moosavi, S.M., Novotny, B.Á., Ongari, D., Moubarak, E., Asgari, M., Kadioglu, Ö., Charalambous, C., Ortega-Guerrero, A., Farmahini, A.H., Sarkisov, L., et al. (2022). A data-science approach to predict the heat capacity of nanoporous materials. *Nat. Mater.* 21, 1419–1425. <https://doi.org/10.1038/s41563-022-01374-3>.
- Schmidt, J., Marques, M.R.G., Botti, S., and Marques, M.A.L. (2019). Recent advances and applications of machine learning in solid-state materials science. *NPJ Comput. Mater.* 5, 83. <https://doi.org/10.1038/s41524-019-0221-0>.
- Tawfik, S.A., Isayev, O., Spencer, M.J.S., and Winkler, D.A. (2020). Predicting thermal properties of crystals using machine learning. *Advanced Theory and Simulations. Adv. Theory Simul.* 3, 1900208. <https://doi.org/10.1002/adts.201900208>.
- 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>.
- Kim, S., Noh, J., Gu, G.H., Aspuru-Guzik, A., and Jung, Y. (2020). Generative adversarial networks for crystal structure prediction. *ACS Cent. Sci.* 6, 1412–1420. <https://doi.org/10.1021/acscentsci.0c00426>.
- Long, T., Fortunato, N.M., Opahle, I., Zhang, Y., Samathrakris, I., Shen, C., Gutfleisch, O., and Zhang, H. (2021). Constrained crystals deep convolutional generative adversarial network for the inverse design of crystal structures. *npj Comput. Mater.* 7, 66. <https://doi.org/10.1038/s41524-021-00526-4>.
- Zhao, Y., Al-Fahdi, M., Hu, M., Siriwardane, E.M.D., Song, Y., Nasiri, A., and Hu, J. (2021b). High-throughput discovery of novel cubic crystal materials using deep generative neural networks. *Adv. Sci.* 8, 2100566. <https://doi.org/10.1002/advs.202100566>.
- Lyngby, P., and Thygesen, K.S. (2022). Data-driven discovery of 2D materials by deep generative models. *NPJ Comput. Mater.* 8, 232. <https://doi.org/10.1038/s41524-022-00923-3>.
- Xie, T., Fu, X., Ganea, O.-E., Barzilay, R., and Jaakkola, T. (2021). Crystal diffusion variational autoencoder for periodic material generation. Preprint at arXiv. <https://doi.org/10.48550/arXiv.2110.06197>.
- Yang, M., Cho, K., Merchant, A., Abbeel, P., Schuurmans, D., Mordatch, I., and Cubuk, E.D. (2023b). Scalable diffusion for materials generation. Preprint at arXiv. <https://doi.org/10.48550/arXiv.2311.09235>.
- Zhu, R., Nong, W., Yamazaki, S., and Hippalgaonkar, K. (2024). WyCryst: Wyckoff Inorganic Crystal Generator Framework.

- Matter. <https://doi.org/10.1016/j.matt.2024.05.042>.
- Yang, L., Zhang, Z., Song, Y., Hong, S., Xu, R., Zhao, Y., Zhang, W., Cui, B., and Yang, M.-H. (2023a). Diffusion models: A comprehensive survey of methods and applications. *ACM Comput. Surv.* 56, 1–39. <https://doi.org/10.1145/3626235>.
 - Zeni, C., Pinsler, R., Zügner, D., Fowler, A., Horton, M., Fu, X., Shysheya, S., Crabbé, J., Sun, L., Smith, J., et al. (2023). Mattergen: a generative model for inorganic materials design preprint at arXiv. <https://doi.org/10.48550/arXiv.2312.03687>.
 - Merchant, A., Batzner, S., Schoenholz, S.S., Aykol, M., Cheon, G., and Cubuk, E.D. (2023). Scaling deep learning for materials discovery. *Nature* 624, 80–85. <https://doi.org/10.1038/s41586-023-06735-9>.
 - Szymanski, N.J., Rendy, B., Fei, Y., Kumar, R.E., He, T., Milsted, D., McDermott, M.J., Gallant, M., Cubuk, E.D., Merchant, A., et al. (2023). An autonomous laboratory for the accelerated synthesis of novel materials. *Nature* 624, 86–91. <https://doi.org/10.1038/s41586-023-06734-w>.
 - Piantadosi, S.T. (2014). Zipf's word frequency law in natural language: A critical review and future directions. *Psychon. Bull. Rev.* 21, 1112–1130. <https://doi.org/10.3758/s13423-014-0585-6>.
 - 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, 6000–6010. <https://doi.org/10.48550/arXiv.1706.03762>.
 - Chan, S., Santoro, A., Lampinen, A., Wang, J., Singh, A., Richemond, P., McClelland, J., and Hill, F. (2022). Data distributional properties drive emergent in-context learning in transformers. *Adv. Neural Inf. Process. Syst.* 35, 18878–18891. <https://doi.org/10.48550/arXiv.2205.05055>.
 - 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. <https://doi.org/10.48550/arXiv.2005.14165>.
 - 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>.
 - Galassi, A., Lippi, M., and Torroni, P. (2021). Attention in Natural Language Processing. *IEEE Trans. Neural Netw. Learn. Syst.* 32, 4291–4308. <https://doi.org/10.1109/TNNLS.2020.3019893>.
 - Niu, Z., Zhong, G., and Yu, H. (2021). A review on the attention mechanism of deep learning. *Neurocomputing* 452, 48–62. <https://doi.org/10.1016/j.neucom.2021.03.091>.
 - Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., Mohamed, A., Levy, O., Stoyanov, V., and Zettlemoyer, L. (2019). Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. Preprint at arXiv. <https://doi.org/10.48550/arXiv.1910.13461>.
 - Liao, Y.-L., Wood, B., Das, A., and Smidt, T. (2023). EquiformerV2: Improved Equivariant Transformer for Scaling to Higher-Degree Representations. Preprint at arXiv. <https://doi.org/10.48550/arXiv.2306.12059>.
 - Christiano, P.F., Leike, J., Brown, T., Martic, M., Legg, S., and Amodei, D. (2017). Deep reinforcement learning from human preferences. *Adv. Neural Inf. Process. Syst.* 30, 4302–4310. <https://doi.org/10.48550/arXiv.1706.03741>.
 - Glawe, H., Sanna, A., Gross, E.K.U., and Marques, M.A.L. (2016). The optimal one dimensional periodic table: a modified Pettifor chemical scale from data mining. *New J. Phys.* 18, 093011. <https://doi.org/10.1088/1367-2630/18/9/093011>.
 - Hautier, G., Fischer, C., Ehlacher, V., Jain, A., and Ceder, G. (2011). Data mined ionic substitutions for the discovery of new compounds. *Inorg. Chem.* 50, 656–663. <https://doi.org/10.1021/ic102031h>.
 - 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, 1. <https://doi.org/10.1063/1.4812323>.
 - Davies, D., Butler, K., Jackson, A., Skelton, J., Morita, K., and Walsh, A. (2019). SMACT: Semiconducting materials by analogy and chemical theory. *J. Open Source Softw.* 4, 1361. <https://doi.org/10.21105/joss.01361>.
 - Blum, V., Rossi, M., Kokott, S., and Scheffler, M. (2022). The FHI-aims Code: All-electron, ab initio materials simulations towards the exascale. Preprint at arXiv. <https://doi.org/10.48550/arXiv.2208.12335>.
 - Gao, Y., Shi, E., Deng, S., Shiring, S.B., Snaider, J.M., Liang, C., Yuan, B., Song, R., Janke, S.M., Liebman-Peláez, A., et al. (2019). Molecular engineering of organic-inorganic hybrid perovskites quantum wells. *Nat. Chem.* 11, 1151–1157. <https://doi.org/10.1038/s41557-019-0354-2>.
 - Park, J.Y., Song, R., Liang, J., Jin, L., Wang, K., Li, S., Shi, E., Gao, Y., Zeller, M., Teat, S.J., et al. (2023). Thickness control of organic semiconductor-incorporated perovskites. *Nat. Chem.* 15, 1745–1753. <https://doi.org/10.1038/s41557-023-01311-0>.
 - Pederson, R., Kozłowski, J., Song, R., Beall, J., Ganahl, M., Hauru, M., Lewis, A.G.M., Yao, Y., Mallick, S.B., Blum, V., and Vidal, G. (2023). Large scale quantum chemistry with tensor processing units. *J. Chem. Theory Comput.* 19, 25–32. <https://doi.org/10.1021/acs.jctc.2c00876>.
 - Ratcliff, L.E., Mohr, S., Huhs, G., Deutsch, T., Masella, M., and Genovese, L. (2017). Challenges in large scale quantum mechanical calculations. *WIREs Comput. Mol. Sci.* 7, e1290. <https://doi.org/10.1002/wcms.1290>.
 - Song, R., Liu, C., Kanai, Y., Mitzl, D.B., and Blum, V. (2023). Structure and electronic tunability of acene alkylamine based layered hybrid organic-inorganic perovskites from first principles. *Phys. Rev. Mater.* 7, 084601. <https://doi.org/10.1103/PhysRevMaterials.7.084601>.
 - Cramer, C.J., and Truhlar, D.G. (2009). Density functional theory for transition metals and transition metal chemistry. *Phys. Chem. Chem. Phys.* 11, 10757–10816. <https://doi.org/10.1039/B907148B>.
 - Wang, T., McWhorter, T.M., McKeown Wessler, G.C., Yao, Y., Song, R., Mitzl, D.B., and Blum, V. (2023). Exploration, Prediction, and Experimental Verification of Structure and Optoelectronic Properties in I2-Eu-IV-X4 (I= Li, Cu, Ag; IV= Si, Ge, Sn; X= S, Se) Chalcogenide Semiconductors. *Chem. Mater.* 36, 340–357. <https://doi.org/10.1021/acs.chemmater.3c02218>.
 - Xue, J., Song, R., Guo, Z., Sung, H.H.Y., Lortz, R., Williams, I.D., and Lu, H. (2023). Extended Honeycomb Metal Chloride with Tunable Antiferromagnetic Correlations. *Chem. Mater.* 36, 551–560. <https://doi.org/10.1021/acs.chemmater.3c02681>.
 - Blum, V., Gehrke, R., Hanke, F., Havu, P., Havu, V., Ren, X., Reuter, K., and Scheffler, M. (2009). Ab initio molecular simulations with numeric atom-centered orbitals. *Comput. Phys. Commun.* 180, 2175–2196. <https://doi.org/10.1016/j.cpc.2009.06.022>.
 - Perdew, J.P., Burke, K., and Ernzerhof, M. (1996). Generalized gradient approximation made simple. *Phys. Rev. Lett.* 77, 3865–3868. <https://doi.org/10.1016/j.cpc.2009.08.008>.
 - Huhn, W.P., and Blum, V. (2017). One-hundred-three compound band-structure benchmark of post-self-consistent spin-orbit coupling treatments in density functional theory. *Phys. Rev. Mater.* 1, 033803. <https://doi.org/10.1103/PhysRevMaterials.1.033803>.
 - Zhao, R., Yu, V.W.-z., Zhang, K., Xiao, Y., Zhang, Y., and Blum, V. (2021a). Quasi-four-component method with numeric atom-centered orbitals for relativistic density functional simulations of molecules and solids. *Phys. Rev. B* 103, 245144. <https://doi.org/10.1103/PhysRevB.103.245144>.
 - Havu, V., Blum, V., Havu, P., and Scheffler, M. (2009). Efficient O(N) integration for all-electron electronic structure calculation using numeric basis functions. *J. Comput. Phys.* 228, 8367–8379.
 - Li, J., Liu, J., Baronett, S.A., Liu, M., Wang, L., Li, R., Chen, Y., Li, D., Zhu, Q., and Chen, X.-Q. (2021). Computation and data driven discovery of topological phononic materials. *Nat. Commun.* 12, 1204. <https://doi.org/10.1038/s41467-021-21293-2>.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Deposited data		
The Materials Project	Jain, Anubhav et al. "Commentary: The Materials Project: A materials genome approach to accelerating materials innovation." <i>APL materials</i> 1.1 (2013).	https://doi.org/10.1063/1.4812323
Software and algorithms		
e3nn: a modular PyTorch framework for Euclidean neural networks	Geiger, Mario, and Tess Smidt. "e3nn: Euclidean neural networks." arXiv preprint arXiv:2207.09453 (2022).	https://doi.org/10.48550/arXiv.2207.09453
EquiformerV2	Liao, Yi-Lun, et al. "Equiformerv2: Improved equivariant transformer for scaling to higher-degree representations." International Conference on Learning Representations (2024).	https://doi.org/10.48550/arXiv.2306.12059 ; https://github.com/atomicarchitects/equiformer_v2
Crystal Generative Models	This paper; original code for reported results.	https://github.com/fangzel/CrystalGenerativeModels

RESOURCE AVAILABILITY

Lead contact

Further information and requests for resources should be directed to the lead contact, Dr. Zhantao Chen (zhantao@stanford.edu).

Materials availability

This study did not generate new unique reagents.

Data and code availability

- This paper analyzes existing, publicly available data. The accession information for these datasets are listed in the [key resources table](#).
- All original code has been deposited at the GitHub repository (<https://github.com/fangzel/CrystalGenerativeModels>) and is publicly available.
- All additional information required to reanalyze the data reported in this paper is available from the [lead contact](#) upon request.

METHOD SUMMARY

Model setup

We represent the atomic numbers of species of N atoms in a unit cell of real crystal structures as $Z \in \mathbb{R}^{N \times C}$, where Z is a set of one-hot vectors with length $C = 118$ for N atoms in the unit cell. The perturbations applied to the atomic positions in $D = 3$ dimensional space are denoted as $P \in \mathbb{R}^{N \times D}$. Consequently, our model's output regarding atomic species is described as $\tilde{Z} \in \mathbb{R}^{N \times C}$, the type-0 output of the equivariant graph neural network with C output channels, and the output for positional perturbations is $\tilde{P} \in \mathbb{R}^{N \times D}$, the type-1 output of the equivariant graph neural network.

For pre-training, we train EquiformerV2²⁹ with 12 Transformer blocks, 8 attention heads, 118 output channels, a maximum degree of 6 and a maximum order of 2. We use the Adam optimizer with a learning rate of 1×10^{-4} . Batch sizes increases throughout the training process (see details in [STAR methods](#)).

Loss functions

For the pre-train model, we choose negative log likelihood loss (NLLLoss) applied after LogSoftmax as the loss function for atomic species, and MSE for atomic positions:

$$\begin{aligned}\mathcal{L}_{\text{pretrain}} &= \mathcal{L}_{\text{species}} + w_{\text{position}} \times \mathcal{L}_{\text{position}}, \\ \mathcal{L}_{\text{species}} &= - \sum_{i=1}^N \sum_{j=1}^C Z_{ij} \text{LogSoftmax}(\tilde{Z}_i)_j, \\ \mathcal{L}_{\text{position}} &= \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^D (P_{ij} - \tilde{P}_{ij})^2.\end{aligned}\tag{Equation 1}$$

For fine-tuning regression tasks, we use MSE as the loss function. For fine-tuning classification tasks, the loss function is NLLLoss applied after LogSoftmax.

For GAN, having found that the discriminator tends to outperform the generator rapidly while using cross entropy as the generator loss, and leads to gradient vanishing in our tasks, we use Wasserstein distance $\mathcal{L}_W = \frac{1}{M} \sum_{m=1}^M \sum_{k=1}^K |y_k^{(m)} - \tilde{y}_k^{(m)}|$ as the training loss. Here, M denotes the number of samples in each batch and $K = 2$ for binary classification. The Wasserstein loss for the generator $\hat{G} : \mathbb{R}^{N \times (C+D)} \rightarrow \mathbb{R}^{N \times (C+D)}$ is formulated as

$$\mathcal{L}_W^G = \frac{1}{M} \sum_{m=1}^M \sum_{k=1}^2 |y_k^{(m)} - \text{LogSoftmax}(\hat{D}(\hat{G}(X^{(m)})))_k|,\tag{Equation 2}$$

with $y^{(m)} = (0, 1)$ indicating the label vector for each sample m . The input of the generate $X^{(m)}$ consists of masked atomic species $\tilde{Z}^{(m)} \in \mathbb{R}^{N \times C}$ and perturbed positions $\tilde{P}^{(m)} \in \mathbb{R}^{N \times D}$ augmented based on the m -th crystal structure.

On the other hand, for the Wasserstein loss in the discriminator $\hat{D} : \mathbb{R}^{N \times (C+D)} \rightarrow \mathbb{R}^2$, $y^{(m)}$ is set to $(1, 0)$ for real structure inputs, i.e., $\tilde{y}^{(m)} = \text{LogSoftmax}(\hat{D}(X^{(m)}))$; $y^{(m)} = (0, 1)$ for inputs from the generated structures, i.e., $\tilde{y}^{(m)} = \text{LogSoftmax}(\hat{D}(\hat{G}(X^{(m)})))$. When training the generator, we keep the parameters of the discriminator frozen, and vice versa.

To prevent the generator from producing features specifically tailored to the weakness of the discriminator, rather than learning to produce the realistic crystal structures, we incorporate $\mathcal{L}_{\text{species}}$ and $\mathcal{L}_{\text{position}}$, as specified Equation 6, into the generator loss:

$$\mathcal{L}^G = w_W \times \mathcal{L}_W^G + \mathcal{L}_{\text{species}} + w_{\text{position}} \times \mathcal{L}_{\text{position}}.\tag{Equation 3}$$

Note that we use $w_{\text{position}} = 200$ as for pre-training and GAN. The training of GAN is done with $w_W = 0.1$, though both $w_W = 1$ and 0.1 have been tested with no significant difference observed.

Modifications to the training procedure are also made to avoid this issue, such as reducing the learning rate of the discriminator ($\text{lr}^{G=10^{-4}}$ and $\text{lr}^{D=10^{-5}}$), introducing infinitesimal positional noise to the real crystal structures (the noise follows a normal distribution with standard deviation of $\sigma = 0.001 \times \min(\text{edges})$), and adjusting the training schedule to update the generator ten times more frequent than the discriminator.

Related metrics

Compositional similarity is defined as the average percentage match between the generated and original material compositions across the test set:

$$\text{Comp. Similarity} = \frac{1}{|\text{test set}|} \sum_{m=1}^{|\text{test set}|} \delta(Z^{(m)} - \tilde{Z}^{(m)}),\tag{Equation 4}$$

Here $\tilde{Z}_i^{(m)} = \delta(j - \arg\max_k (\text{LogSoftmax}(\tilde{Z}_i^{(m)})_k))$, where $\tilde{Z}^{(m)} \in \mathbb{R}^{N \times C}$ denotes the compositional output from the generative model, as shown in Equation 6. This metric provides insight into the model's ability to accurately restore the composition of materials. However, it is important to note that (1) while high compositional similarity indicates model accuracy, it is contradictory to assessing its capability to innovate novel structures; (2) for computational simplicity, a generated structure is considered compositionally identical to the original structure only if $Z_i = \tilde{Z}_i$ for all i . This definition, however, overlooks the scenario where two atoms of different species ($Z_i \neq \tilde{Z}_j$) may have their locations exchanged due to the large positional noise. Such an exchange would lead the generative models to swap their species and positions, resulting in $\tilde{Z}_i = Z_j$, $\tilde{Z}_j = Z_i$, $\tilde{P}_i = P_j$, and $\tilde{P}_j = P_i$. In this case, the generated structure should still be considered equivalent to the original counterpart. With low noise levels, the probability of this occurrence is sufficiently low. However, as our analysis in Figure 4B includes high noise levels, this could lead to inaccuracies in assessing compositional similarity.

Structural similarity, on the other hand, aims to evaluate the model's ability in replicating the structural integrity when the model decides to generate compositions identical to the originals. It is defined as:

$$\text{Struc. Similarity} = 1 - \frac{1}{|\text{test set}|} \sum_{m=1}^{|\text{test set}|} \delta(Z^{(m)} - \tilde{Z}^{(m)}) \left(\frac{1}{N_m} \sum_{i=1}^{N_m} \sqrt{\sum_{j=1}^D (P_{ij}^{(m)} - \tilde{P}_{ij}^{(m)})^2} \right) / S^{(m)}.\tag{Equation 5}$$

Here, the equation considers the Euclidean distance between the positions in the original structure ($P^{(m)} \in \mathbb{R}^{N \times D}$) and generated structure ($\tilde{P}^{(m)} \in \mathbb{R}^{N \times D}$) for the m -th test sample, normalized by the average positional deviation from the equilibrium introduced to the input ($S^{(m)} \in \mathbb{R}$). A higher structural similarity score, approaching 1, indicates a closer match to the original material's structural positioning, offering a direct measure of the model performance in restoring material structure.

For DFT calculations, we employ the all-electron electronic structure package FHI-AIMS⁴⁴ to conduct single-point DFT calculations. We specifically exclude systems containing partially-filled d - or f -orbitals in this part, so the combination of semilocal Perdew-Burke-Ernzerhof (PBE) exchange-correlation (XC) functional⁴⁵ and the “atomic ZORA” scalar relativity^{44,46,47} is considered to be suitable for total energy prediction. Moreover, “intermediate” numerical settings of the numeric atom-centered-orbital basis are used,⁴⁸ and the k -grids for different calculations are adjusted according to different unit cell sizes to ensure adequate density of the sampled points in the reciprocal space.

Motivations and discussions

In this study, we propose a new framework for autonomous material generation. The distinctive feature of our approach is the minimal reliance on pre-existing human pre-knowledge during the training process: the generative model is trained exclusively with the atomic numbers of species and their positions for each material structure. This approach not only significantly reduces computational costs and mitigates biases arising from prior assumptions, but also facilitates the study of mappings from material structures to complex physical properties.

Motivation

Our work draws inspiration from the field of NLP, which has seen remarkable advancements in recent years. These improvements have been largely driven by the adoption of self-supervised learning strategies²⁵ and, though more resource-intensive, in-context learning,²⁴ as well as the self-attention mechanism in LLMs.^{22,26,27} Besides these aspects, the data distribution has been identified as a crucial factor in driving emergent behaviors in transformer-based LLMs.²³ As shown in Figure S1A, we observe that the distribution of elemental abundance in crystal structures is close to that of vocabulary abundance in natural languages, following a power law distribution—a common pattern in natural datasets like languages. This observation suggests that the success of LLMs could be beneficial in material prediction research. Therefore, we trained a generative model for crystal structures in a similar manner, as detailed in Section.

Obstacles and detours

The challenge, however, lies in the absence of a critic to evaluate the generated structures. We propose utilizing the actor-critic learning framework to guide generative training, similar to applying RLHF in NLP. Ideally, numerical methods like DFT calculations or even human evaluations could serve as the critic, while employing GANs offers a computationally efficient solution for early-stage training. More importantly, our objectives are not limited to merely aiding in material generation; we are also keen on exploring the insights that can be extracted from crystal structures. To ensure our training process remain unbiased and uninfluenced by pre-existing knowledge, we deliberately avoided using DFT and similar methods during the training phase. As a result, we choose to integrate the GAN architecture into our training pipeline to fully exploit existing data without relying on external sources like DFT-calculated stability labels (refer to Section).

Solution

To autonomously generate physically meaningful crystals, the model is expected to self-learn the principles of crystal formation and potentially gaining inferential capabilities. The training of material prediction can be mainly divided into two stages: (1) *Pre-training stage*: we guide the model through fill-in-the-blank and sentence-arranging exercises, using known crystal structures as reference. (2) *Generative adversarial stage*: two sufficient trained models play opposing roles. One model endeavors to generate structures indistinguishable from real ones, while the other one attempts to distinguish between the generated and real ones. Beyond generative learning, our model also demonstrate the adaptability for down-stream regression and classification tasks through supervised fine-tuning (see Section).

Generation or reconstruction?

Our current model focuses on generating the optimal crystal structure under given crystallographic design—it may not yield the globally optimal structure if the specified conditions do not fall within the convex of global optima. It worth emphasizing that our models has never seen and been trained by the test samples. Therefore, if it performs well in “reconstructing” a never-before-seen corrupted inputs, we would also expect it to perform well with any input—without the need to resemble any existing material structure. We demonstrate the models’ reconstruction ability in Section.

Comparison

The prevailing popular models for crystal structure generation utilize diffusion process and VAEs. Diffusion models are typically applied to generate crystals from given composition, unit cell information, and null position embeddings. Our training method asymptotically approaches to a diffusion process as the noise introduced to the inputs goes to a large limit (requiring an additional modification to the loss function). Currently, our generative model is presently confined to scenarios where the basic design of the desired

crystal structure is predefined. In other words, while the diffusion-based studies target the global optimal structures for certain compositions, our work focuses on identifying local optima within a defined framework. Nonetheless, our training strategy can be adapted for global optimization by introducing higher noise levels. VEAs follow a distinct approach: projecting crystal structures into a latent space through an encoder and then decoding it back to crystal structures, enabling material generation through sampling over the latent space.

Probabilistic model

Our generative model functions as a conditional probabilistic model, as it outputs the probability of specific atomic species at each masked site. Therefore, this model offers a probabilistic understanding into the nature of crystals in an interpretable manner, delivering richer information compared to other unconditional probabilistic models.^{31,32} Details are provided in Section.

Self-supervised pre-training stage

Model setup

Our study utilizes the most recent equivariant graph neural network with Transformer architecture, EquiformerV2,²⁹ for our pre-trained model. Equivariance is a property of an operator $\Phi : X \rightarrow Y$ mapping between vector spaces X and Y by which it commutes with the group action of a group (G, \circ) , i.e., $\Phi \circ \rho^X(g) = \rho^Y(g) \circ \Phi$, where $\rho^X(g)$ is the group representation ρ of $g \in G$ action on X . For our purposes, the equivariant graph neural network acts as the operator Φ , being equivariant to $E(3)$ group actions—3D translations, rotations and inversions—in processing 3D atomistic structures X and Y . This allows the model to inherently capture symmetries in the data, eliminating the need for data augmentation. During the pre-training, we trained EquiformerV2 with 12 Transformer blocks, 8 attention heads, 118 output channels, a maximum degree of 6 and a maximum order of 2. We used the Adam optimizer with a learning rate of 1×10^{-4} . Batch sizes were adjusted through the training process: 1 for early epochs (below 25), 3 for mid-stage epochs (25–45), and 5 for later epochs (beyond 45).

Dataset

We collect a dataset comprising 21,250 diverse, stable crystal structures from the Materials Project,³³ allocating 80% for training and 10% each for validation and testing. Our selection criteria prioritize samples with available electronic band structures and DOS to enhance the quality and reliability of the training data. The pre-train model receives input crystal structures undergoing augmentation through two operations: (1) with a 50% probability, randomly selecting a species and masking all its atoms, and with a 50% probability, masking 15% of the atoms within a unit cell; (2) introducing random perturbations to the positions of all atoms. The perturbations follow a Gaussian distribution with a standard deviation of $\tilde{\sigma}_{\text{noise}} \times \min(\text{edges})$, where $\tilde{\sigma}_{\text{noise}} = 0.1$ is set for the pre-training phase. Our model is designed to reconstruct these corrupted crystal structure by revealing the masked atoms and accurately adjusting the positions of all atoms.

In our representation, the atomic species of N atoms within a unit cell of a real crystal structure are represented as $Z \in \mathbb{R}^{N \times C}$. Here, Z constitutes a set of one-hot vectors, each of length $C = 118$, corresponding to the N atoms in a unit cell. The perturbations applied to the real atomic positions in $D = 3$ dimensional space are denoted as $P \in \mathbb{R}^{N \times D}$. Consequently, the output of our model for predicting atomic species is denoted as $\tilde{Z} \in \mathbb{R}^{N \times C}$, which is the type-0 output of the equivariant neural network with C output channels, and the output for positional perturbations is $\tilde{P} \in \mathbb{R}^{N \times D}$, which is the type-1 output of the equivariant graph neural network.

Pre-training losses

To optimize the model weights for this task, the training loss for each sample is defined as

$$\begin{aligned} \mathcal{L}_{\text{pretrain}} &= \mathcal{L}_{\text{species}} + w_{\text{pos}} \times \mathcal{L}_{\text{pos}}, \\ \mathcal{L}_{\text{species}} &= - \sum_{i=1}^N \sum_{j=1}^C Z_{ij} \text{LogSoftmax}(\tilde{Z}_i)_j = - \sum_{i=1}^N \sum_{j=1}^C Z_{ij} \log \frac{\exp \tilde{Z}_{ij}}{\sum_{k=1}^C \exp \tilde{Z}_{ik}}, \\ \mathcal{L}_{\text{pos}} &= \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^D (P_{ij} - \tilde{P}_{ij})^2. \end{aligned} \tag{Equation 6}$$

In Figure S2, we investigated the impact of the weight w_{pos} during pre-training. $\mathcal{L}_{\text{species}}$ remains unaffected by changes in w_{pos} , whereas \mathcal{L}_{pos} decreases with an increase in w_{pos} and reaches a plateau when $w_{\text{pos}} \approx 200$. Based on this observations, we selected $w_{\text{pos}} = 200$ for our pre-training. The pre-training losses, normalized by the number of atoms in each batch, is shown in Figure S3.

Supervised fine-tuning stage

Strategy

Through pre-training on a broad and diverse dataset, the model develops a fundamental understanding of the principles of materials composition and encodes this information in the model parameters. The pre-trained model can be more quickly and efficiently tailored for a variety of downstream tasks than a model trained from scratch. For a specific supervised task, the pre-trained model is merged with a feedforward

layer, which is randomly initialized and is designed to adapt the output of the pre-trained model to the requirement of the new task. During fine-tuning, both the pre-trained models and the feedforward layer are trained on a task-specific dataset.

Fine-tuning on regression tasks

For regression tasks, we evaluate our model on predicting (1) the Fermi energy, (2) bulk moduli and shear moduli, and (3) phonon density-of-states (DOS) based solely on the crystal structure, using MSE as the loss function. (1) For the prediction of Fermi energy, we utilize the same training set as used in the pre-training phase to prevent any information from the validation set being inadvertently exposed to the model. We further explore the influence of the training set size, depicted in Figure S4, reveals that the validation loss is significantly affected by the size of train data. (2) To investigate the behavior of fine-tuning on smaller datasets, we fine-tune the model for elasticity properties like bulk moduli and shear moduli with units of logGPa. This fine-tuning is performed using a dataset comprising 1915 training samples and 638 validation samples from Materials Project. (3) The pre-trained model is originally trained by predicting atomic displacement in crystal structures, which may capture how local structural perturbations influence the vibration properties of materials, and establishing a direct correlation between structural changes and macroscopic phononic behavior. Therefore, we leverage our pre-trained model to predict the phonon DOS, as shown in Figure S4, with 1,524 training samples and 304 validation samples from.⁴⁹

Fine-tuning on classification tasks

For classification tasks, the loss function is Negative Log Likelihood Loss (NLLLoss) applied after LogSoftmax, as $\mathcal{L}_{\text{species}}$ in Equation 6. To assess our pre-trained model's capability in classifying crystal structures, we naturally focus on identifying the stability of crystal structures. We fine-tune the model by using 4,017 stable and 4,065 unstable cubic crystal structures from Materials Project. It is worth noting this model was initially pre-trained exclusively on stable crystal structures. For evaluation, we utilize 2,000 samples from each crystal types, except for triclinic crystal structures, which have only 1,483 samples. These samples are nearly evenly split between stable and unstable structures. The accuracy of stability prediction is shown in Figure S5. Remarkably, although the model's fine-tuning was confined to cubic crystals, it demonstrates extrapolative capabilities to other types of crystal structures.

Generative adversarial network stage

Motivation

The objective of our generative model extends beyond merely reconstruct the original crystal structure from its contaminated counterpart, while also exploring novel crystal structures that have not been discovered before. The challenge, however, lies in the limitation of the loss function, which can only optimize the reconstruction goals, and the absence of a tool for discriminating generated structures.

We consider the framework of actor-critic learning to guide the generative training. Based on this framework, we have two options: self-supervised GANs or incorporating human pre-knowledge of crystal structures as a critic. Compared to GANs, the later one, such as utilizing DFT, are computationally expensive, with duration ranging from a few minutes to several hours for each crystal graph. Furthermore, the critic neural network model supervised by substantial DFT data has been sufficiently trained in previous studies.¹⁹ Such well-trained critic model could serve as a valuable tool for the future enhancement of our generative model.

Another concern involves utilizing generative training to explore how much information is embedded in the crystal structures. This topic is analogous to examining whether a child, without any formal teaching on recognizing sentiments, can discern the sentimental tones in a sentence after extensive reading. To ensure rigorous control over variables, we tend to isolate the teaching from humans.

Model setup

For the above reasons, we finally select self-supervised GANs. In this setup, the discriminative neural network is trained to distinguish the crystal structures produced by the generative model from the real crystal structures. Concurrently, the generative neural network is optimized to deceive the discriminator. We initiated the generator with pre-trained model parameters and used the fine-tuning model parameters for stability classification in downstream tasks as the discriminator's starting point, which is chosen because this task shares some commonalities with the discriminator's role, despite their tasks not being entirely aligned.

Generative adversarial network training losses

Having found that the discriminator tends to outperform generator rapidly while using cross entropy as the generator loss, and leads to gradient vanishing in our tasks, we use Wasserstein distance as an alternative:

$$\mathcal{L}_{\text{Wasserstein}} = \frac{1}{M} \sum_{m=1}^M \sum_{k=1}^K \left| y_k^{(m)} - \tilde{y}_k^{(m)} \right|,$$

where M is the number of samples in each batch and $K = 2$ for binary classification. The Wasserstein loss for the generator $\hat{G} : \mathbb{R}^{N \times (C+D)} \rightarrow \mathbb{R}^{N \times (C+D)}$ is formulated as

$$\mathcal{L}_{\text{Wasserstein}}^G = \frac{1}{M} \sum_{m=1}^M \sum_{k=1}^2 \left| y_k^{(m)} - \text{LogSoftmax} \left(\widehat{D} \left(\widehat{G} \left(X^{(m)} \right) \right) \right)_k \right|,$$

with $y^{(m)} = (0, 1)$ indicating the label vector for each sample m . The input of the generator, $X^{(m)}$, consists of masked atomic species $\tilde{Z}^{(m)} \in \mathbb{R}^{N \times C}$ and perturbed positions $\tilde{P}^{(m)} \in \mathbb{R}^{N \times D}$ augmented based on the m -th crystal structure.

On the other hand, for the Wasserstein loss in the discriminator $\widehat{D} : \mathbb{R}^{N \times (C+D)} \rightarrow \mathbb{R}^2$, $y^{(m)}$ is set to $(1, 0)$ for real structure inputs, i.e., $\tilde{y}^{(m)} = \text{LogSoftmax}(\widehat{D}(X^{(m)}))$; $y^{(m)} = (0, 1)$ for inputs from the generated structures, i.e., $\tilde{y}^{(m)} = \text{LogSoftmax}(\widehat{D}(\widehat{G}(X^{(m)})))$. When training the generator, we keep the parameters of the discriminator frozen, and vice versa.

To prevent the generator from producing features specifically tailored to the weakness of the discriminator, rather than learning to produce the realistic crystal structures, we incorporate $\mathcal{L}_{\text{species}}$ and \mathcal{L}_{pos} , as specified Equation 6, into the generator loss:

$$\mathcal{L}^G = w_W \times \mathcal{L}_{\text{Wasserstein}}^G + \mathcal{L}_{\text{species}} + w_{\text{pos}} \times \mathcal{L}_{\text{pos}}.$$

Note that we use $w_{\text{pos}} = 200$ as for pre-training and GAN. The training of GAN is done with $w_W = 0.1$, though both $w_W = 1$ and 0.1 have been tested with no significant difference observed.

Modifications to the training procedure are also made to alleviate the issue of the discriminator outperforming the generator, such as reducing the learning rate of the discriminator ($\text{lr}_G = 10^{-4}$ and $\text{lr}_D = 10^{-5}$), positional noise adding to the real crystal structures (the noise follows a normal distribution with a standard deviation of $0.001 \times \min(\text{edges})$), and adjusting the training schedule to update the generator more frequently than the discriminator.

Results

In the generative adversarial stage, both the training and validation datasets employed are identical to those used in the pre-training, and the same level of positional noise is introduced to the input structures, i.e., $\sigma_{\text{noise}} = 0.1 \times \min(\text{edges})$, as shown in Figure S6. Despite observing minimal reduction in the loss $\mathcal{L}_{\text{species}}$ and \mathcal{L}_{pos} across increasing training epochs, the GAN achieves a better performance in generating physically reasonable crystal structures, at least visually, compared to the pre-train model.

To make our model compatible with larger displacements, we experimented positional noise with $\sigma_{\text{noise}} = 0.2 \times \min(\text{edges})$ during GAN training in Figure S7. This adjustment led to notable decreases in $\mathcal{L}_{\text{species}}$ and \mathcal{L}_{pos} . We also evaluated the performance of the pre-train model (trained with noise level $\sigma_{\text{noise}} = 0.1 \times \min(\text{edges})$) and the generator (trained with noise level $\sigma_{\text{noise}} = 0.2 \times \min(\text{edges})$) on reconstructing structures perturbed with noise at $\sigma_{\text{noise}} = 0.2 \times \min(\text{edges})$. As illustrated in Figure S7C, the pre-train model shows a decline in predicting atomic species, while the GAN generator exhibited an improved performance in handling the structures with a higher noise level.

In addition to the main text, we provide additional visual illustrations to demonstrate the generative capabilities of our models, as shown in Figure S8.

Comparative analysis of generative capability

Validity

To quantitatively evaluate the crystal structures generated by our generative models, we present analysis of *structural validity*, which depends on distance between atoms, and *composition validity*, based on the overall charge.¹⁴ The results are shown in the main text. Here, we further analyze the generated structures that do not pass the validity tests by analyzing the number of atoms per unit cell and the elements in the invalid structures in Figures S9 and S10. These figures infer the following arguments.

- (1) When the noise level in the testing inputs matches that of the training inputs, the GAN generator, though not outperforming the pre-trained model in constructing structurally valid crystals, demonstrates enhanced robustness in generating compositionally valid structures. Despite all generative models being trained on inputs with only 15% of atoms masked, the GAN approach improves extrapolation capability, especially in scenarios where 30% atoms are randomly masked (mask type γ).
- (2) When the models are tested with a higher noise level, as depicted in Figures S9B and S10B, the GAN generators surpasses the pre-trained in creating compositionally and structurally valid crystals, especially obtaining a better performance when the number of atoms is large.
- (3) Interestingly, the GAN generator trained with a higher noise level ($G_{0.2}$) fails to achieve the highest validity on the the test set with a lower noise level ($\bar{\sigma}_{\text{noise}} = 0.1$). This observation suggests that diversifying the test set by including variations in $\bar{\sigma}_{\text{noise}}$ may enhance the model's versatility.
- (4) Figure S10B depicts the relationship between the crystal structures that models fail to successfully generate and the number of atoms (per unit cell) within these structures. The similarity in the distribution of the number of atoms in compositionally invalid structures to that in original structures (Figure S1B) suggests a weak correlation between composition validity and the number of atoms. In contrast, the distribution for structure invalidity deviates from that observed in realistic structures. Although it is not perfectly reliable to draw statistically significant conclusions for rigorous analysis from the structure invalidity distribution, due to the relatively small number (approximately 10 per set) of invalid structures generated by our models in each augmented dataset, the trends observed still demonstrate the GAN generator's enhanced capability in structure validity with increased $e\bar{\sigma}_{\text{noise}}$.

- (5) The pattern of invalidity distribution provides insights for future model improvements, for example, by training with crystal structures that are similar to those fall within the lower validity ranges.

Similarity

We introduce two metrics designed to evaluate the effectiveness of generative models in accurately reproducing the properties of ground truth materials. These metrics, namely compositional similarity and structural similarity, quantify the accuracy with which the models replicate material structures in alignment with their original counterparts.

Compositional similarity is defined as the average percentage match between the generated and original material compositions across the test set:

$$\text{Compositional Similarity} = \frac{1}{|\text{test set}|} \sum_{m=1}^{|\text{test set}|} \delta(Z^{(m)} - \tilde{Z}^{(m)}),$$

Here $\tilde{Z}_j^{(m)} = \delta(j - \arg\max_k (\text{LogSoftmax}(\tilde{Z}_i^{(m)})_k))$, where $\tilde{Z}^{(m)} \in \mathbb{R}^{N \times C}$ denotes the compositional output from the generative model, as shown in Equation 6. This metric provides insight into the model's ability to accurately restore the composition of materials. However, it is important to note that while high compositional similarity indicates model accuracy, it is contradictory to assessing its capability to innovate novel structures.

It is important to clarify that, for computational simplicity, a generated structure is considered compositionally identical to the original structure only if $Z_i = \tilde{Z}_i$ for all i . This definition, however, overlooks the scenario where two atoms of different species ($Z_i \neq \tilde{Z}_i$) may have their locations exchanged due to the large positional noise. Such an exchange would lead the generative models to swap their species and positions, resulting in $\tilde{Z}_i = Z_j$, $\tilde{Z}_j = Z_i$, $\tilde{P}_i = P_j$, and $\tilde{P}_j = P_i$. In this case, the generated structure should still be considered equivalent to the original counterpart. When noise level is sufficiently low, the likelihood of this scenario occurring is minimal, but as our analysis in Figure S11 includes large noise levels, it is conceivable that these conditions may induce inaccuracies in assessing compositional similarity.

Structural similarity, on the other hand, aims to evaluate the model's ability in replicating the structural integrity when the model decides to generate compositions identical to the originals. It is defined as:

$$\text{Structural Similarity} = 1 - \frac{1}{|\text{test set}|} \sum_{m=1}^{|\text{test set}|} \delta(Z^{(m)} - \tilde{Z}^{(m)}) \left(\frac{1}{N_m} \sum_{i=1}^{N_m} \sqrt{\sum_{j=1}^D (P_{ij}^{(m)} - \tilde{P}_{ij}^{(m)})^2} \right) / S^{(m)}.$$

Here, the equation considers the Euclidean distance between the positions in the original structure ($P^{(m)} \in \mathbb{R}^{N \times D}$) and generated structure ($\tilde{P}^{(m)} \in \mathbb{R}^{N \times D}$) for the m -th test sample, normalized by the average positional deviation introduced to the input ($S^{(m)} \in \mathbb{R}$). A higher structural similarity score, approaching 1, indicates a closer match to the original material's structural positioning, offering a direct measure of the model performance in restoring material structure.

To facilitate an evaluation of model performance on various corrupted test samples, we illustrate the distribution of the compositional and structural similarity across noise levels in Figure S11A. We further conduct a comparative analysis of the performances of three generative models within the same test dataset in Figure S11B. From these analyses, we can infer the following key insights.

- (1) The distribution of compositional and structural similarities approximates a Gaussian distribution, with the mean values locating at the standard deviation of the noise distribution introduced to the training samples, $\sigma_{\text{noise}} = \bar{\sigma}_{\text{noise}} \times \min(\text{edge})$.
- (2) GAN generators perform better across all test sets when compared to the pre-trained model eG_{pr} .
- (3) When handling the test samples with mask type β (randomly selecting and masking 15% of atoms), generative models are more inclined to replicate the original crystals. When the model has decided to restore the originals, their performances on restoring the atoms back to the original positions are almost unaffected by the specific mask types used in the test samples.
- (4) As the noise level of the test set increases to $\bar{\sigma}_{\text{noise}} = 0.4$, the models quickly fail to revert the compositions and structures back to their originals. Our model serves as a mechanism to find the local minima in incomplete crystal structure, rather than finding a global minimum. At sufficiently high noise levels, the input structure is pushed into the other local minima, thus making the model unlikely to reconstruct a structure that matches the ground truth exactly.
- (5) $G_{0.2}$ uses a training set with $\bar{\sigma}_{\text{noise}} = 0.2$ during GAN training, but since it utilizes the model parameters of the pre-trained model G_{pre} , hence it has been effectively exposed to training set with different noise levels. This exposure helps the model perform better on training sets with higher noise levels without losing performance on the training set at the lower noise levels. It suggests that sampling the standard deviation of noise distribution of the training set can enable the model to handle test samples with different noise levels, potentially improving its accuracy in finding both local and global minima in the crystal structures.

Figures S12 and S13 visualize the effects of the size of crystals and the species that the models should unmask on models' performance. The insights drawn from these figures are summarized as follows.

- (1) The generative models demonstrate increased difficulty in replicating the compositions of the crystals that have a large number of atoms and contain species that are less frequent in the test set.

- (2) Despite the non-uniform distribution of the size of crystals and species abundance within the datasets (as shown in Figure S1), the structural similarity scores for all models are marginally affected by the number of atoms and the types of species masked at low noise levels.
- (3) At high noise levels, the improvements of $\mathcal{G}_{0.2}$ are evident in its enhanced capability in reconstructing crystal structures compared to other generative models, assuming our original crystal structures represent the global optima. This improvement is particularly noticeable when tested with a crystal having a larger number of atoms and less common species being masked.

In Figure S14, we present a summary of the **validity** and **similarity** scores, highlighting the highest score achieved in each test set, in which samples have been masked and perturbed in different ways.

Density functional theory calculations

In Figure S15, we present a comparison of the structural similarity and relative total energy of generated crystal structures. These structures are generated by various generative models with inputting contaminated structures randomly sampled from the test set. We draw the following conclusions.

- (1) The generated structure exhibiting the lowest relative total energy ΔE shows the highest similarity for each crystal. This observation suggests that structural similarity can serve as an indicator of part of the information provided by DFT calculations.
- (2) Structures produced by GAN generators generally demonstrate a lower ΔE and a higher structural similarity, indicating the effectiveness of GANs in reconstructing structures with lower total energy compared to those generated by the pre-trained model.
- (3) Although the generated structures are nearly visually indistinguishable from their original counterparts, their structural similarities are around 50%. This is attributed to the normalization of this metric by the average noise.

Novelty

In this study so far, we have primarily demonstrated the model's capability for reconstruction, but have not yet characterized its potential for generation. Therefore, we define a new quantity termed *novelty*. A generated structure is considered novel if it is both compositionally and structurally valid, and it exhibits a distinct composition from the original. In Figure S17, we present the rate of atomic species being replaced or replacing other species under three different masking types. Here, we denote $n_{\text{replaced}}^{(\mathcal{G},\delta)}(X)$ the number of novel crystals generated by our generative model \mathcal{G} , under mask type (δ) , where species s is replaced, and $n_{\text{replace}}^{(\mathcal{G},\delta)}(X)$ as the counterpart for the species X substituting original species, and $N_{\text{novel}}^{(\mathcal{G},\delta)}$ as the total number of novel crystals generated under mask type (δ) . The rate of species X replacing other species (or being replaced) can be defined as

$$\text{Rate}_{\text{replace}(d)}^{(\mathcal{G},\delta)}(X) = n_{\text{replace}(d)}^{(\mathcal{G},\delta)}(X) / N_{\text{novel}}^{(\mathcal{G},\delta)}. \quad (\text{Equation 7})$$

We notice that simple substitution methods³¹ compute the likelihood of substituting species A with B , denoted as $P(A \rightarrow B)$, through data mining. This method is particularly useful when the probability of one species being replaced by another remains constant, irrespective of different atomic and positional information within various crystals—although such scenarios that rarely applies. On the other hand, our generative model is essentially a more sophisticated probabilistic model, capable of offering conditional probabilities $P(A \rightarrow B | \text{condition})$ under diverse conditions, such as varying structures and compositions. When conditions are overlooked, as shown in Figure S16, the averaged probabilities $\frac{1}{|\text{condition}|} \sum_{\text{condition}} P(A \rightarrow B | \text{condition})$ given by our generative models bears a resemblance to those produced by simple substitution.

In Figure S17A, we expand our analysis by exploring the correlation between the species being replaced and those replacing them in the valid, novel crystal structure, generated by the pre-trained model \mathcal{G}_{pre} and the GAN generator $\mathcal{G}_{0.1}$ (refer to the main text for $\mathcal{G}_{0.2}$), computed over augmented test data for all mask types and positional noise levels. This results, as depicted in the figure, indicate that within each group block, the correlations are predominant. This suggests that our generative models are capable of categorizing elements into groups in a manner analogous to the classification system used in the periodic table.

Figures S17B and S17C presents the detailed distribution of $\text{Rate}_{\text{replace}(d)}^{(\mathcal{G},\delta)}(X)$ across different mask types, noise levels, and models. In this figure, the elements are ordered based on their occurrence probability in our dataset. The relatively uniform distributions implies that the model is not affected by the uneven distribution of species occurrences in the training set, thereby reducing potential bias. When the positional noise in the test input crystals exceeds that of the training inputs, we observe a tendency of the models to generate novel crystals. However, it is worth noting that while these crystals have passed the composition and structural validity tests, rigorous testing is required to prove the stability of their chemical structures.