



Establishing baselines for generative discovery of inorganic crystals†

Nathan J. Szymanski and Christopher J. Bartel *Cite this: *Mater. Horiz.*, 2025, 12, 8000Received 4th January 2025,
Accepted 3rd July 2025

DOI: 10.1039/d5mh00010f

rsc.li/materials-horizons

Generative artificial intelligence offers a promising avenue for materials discovery, yet its advantages over traditional methods remain unclear. In this work, we introduce and benchmark two baseline approaches – random enumeration of charge-balanced prototypes and data-driven ion exchange of known compounds – against four generative techniques based on diffusion models, variational autoencoders, and large language models. Our results show that established methods such as ion exchange are better at generating novel materials that are stable, although many of these closely resemble known compounds. In contrast, generative models excel at proposing novel structural frameworks and, when sufficient training data is available, can more effectively target properties such as electronic band gap and bulk modulus. To enhance the performance of both the baseline and generative approaches, we implement a post-generation screening step in which all proposed structures are passed through stability and property filters from pre-trained machine learning models including universal interatomic potentials. This low-cost filtering step leads to substantial improvement in the success rates of all methods, remains computationally efficient, and ultimately provides a practical pathway toward more effective generative strategies for materials discovery. By establishing baselines for comparison, this work highlights opportunities for continued advancement of generative models, especially for the targeted generation of novel materials that are thermodynamically stable.

Introduction

The discovery of new materials has long been a cornerstone of technological progress, driving many of the innovations that shape modern society.¹ Breakthroughs in layered and Li-rich battery cathodes, for example, have enabled the widespread

University of Minnesota, Department of Chemical Engineering and Materials Science, Minneapolis, MN, USA 55455. E-mail: cbartel@umn.edu

† Electronic supplementary information (ESI) available. See DOI: <https://doi.org/10.1039/d5mh00010f>

New concepts

This work establishes comprehensive baselines for the generative discovery of inorganic crystals, comparing traditional methods with generative approaches. By developing and benchmarking two baselines – random enumeration of charge-balanced prototypes and data-driven ion exchange – against three generative models, we demonstrate a new framework for assessing trade-offs in generative materials discovery. The baseline methods excel at proposing stable materials, while the generative models offer greater structural novelty and outperform baselines in identifying materials with targeted properties, such as band gaps near 3 eV. However, both approaches face challenges in finding materials with exceptional properties, such as maximizing bulk modulus, underscoring the need for more diverse training data. By introducing machine learning-based screening as a post-generation filtering step, we significantly enhance success rates across all approaches. This integration highlights a subtle yet crucial balance between stability, novelty, and property optimization, offering a clear framework to evaluate generative AI in materials science while supporting the advancement of future models.

adoption of portable electronics and electric vehicles.² Transparent conducting oxides such as indium tin oxide (ITO) and indium gallium zinc oxide (IGZO) have been critical for the development of touch screens, solar cells, and flat-panel displays.³ Similarly, the discovery of cuprate superconductors in the 1980s reignited interest in high-temperature superconductivity, which remains the subject of extensive research.⁴ These examples highlight the role of materials discovery in advancing transformative technologies and underscore the need for innovative approaches to accelerate future breakthroughs.

The recent emergence of generative artificial intelligence (AI) offers a promising route for designing new materials, particularly inorganic crystals.^{5,6} Early efforts focused on generative adversarial networks (GANs)^{7–9} and variational autoencoders (VAEs),^{10–14} while more recent developments include large language models (LLMs),^{15–19} diffusion-based techniques,^{20–24} normalizing flows,^{25–27} and geodesic random walks.²⁸ These models are often trained on computed materials from open databases such as the Materials Project²⁹ to generate thermodynamically

stable structures, with some also conditioned on specific properties for application-driven campaigns.

There exists a growing number of successes in generative AI for materials design, with validation provided by *ab initio* calculations and experimental synthesis. For example, text-based models such as Chemeleon³⁰ have leveraged contrastive learning to align crystal structures with natural language descriptions, enabling composition- and structure-conditioned generation in a variety of chemical spaces – most notably achieving successful phase prediction in the Li–P–S–Cl system relevant to solid-state batteries. Other models such as FlowLLM²⁶ combine language representations with Riemannian flow matching to refine generated structures, increasing their stability rate threefold. Among diffusion models, MatterGen²² has emerged as a particularly effective method capable of generating materials with targeted chemistry, symmetry, and functional properties – even leading to the synthesis of an AI-generated compound, TaCr₂O₆, whose experimentally measured bulk modulus was within 20% of the predicted value.

Despite these recent successes, it remains difficult to systematically assess the performance of different generative models in a consistent fashion. Tools like matbench-genmetrics provide important frameworks and metrics for evaluating the validity of structures proposed by generative models,³¹ while matbench-discovery addresses the challenge of benchmarking stability predictions made by machine learning (ML) models and interatomic potentials.³² Yet, the extent to which generative models outperform established methods, such as ion exchange or high-throughput screening, is not yet fully understood. Baselines are therefore essential to clarify where these models offer the greatest advantages – whether in producing stable materials, generating novel structures, or achieving targeted properties – and to identify their limitations. Such benchmarks are key to integrating generative models into existing workflows and driving tangible progress in materials discovery.

In this work, we establish two baseline methods for the generation of inorganic crystals: random enumeration of charge-balanced chemical formulae in structure prototypes sourced from the AFLOW database,^{33,34} and ion exchange performed on stable compounds with desired properties from the Materials Project.²⁹ These methods are benchmarked against four generative models – CrystalLM,¹⁵ FTCP,¹² CDVAE,¹³ and MatterGen²² – for the generation of (1) materials that are stable and novel, (2) materials with a band gap near 3 eV, and (3) materials with high bulk modulus. We also integrate two graph neural networks, CHGNet³⁵ and CGCNN,³⁶ to filter and retain generated materials predicted to be stable or exhibit desired properties. This evaluation sheds light on the comparative strengths and weaknesses of traditional and generative approaches to materials discovery, while also providing a set of baselines against which future generative models can be benchmarked.

Methods

Random enumeration

In this method, we randomly paired a known structure prototype with a set of elements also chosen at random. The

prototypes were drawn from 1783 structures listed in the encyclopedia of crystallographic prototypes (AFLOW).^{33,34} Compositional sets were created using three to five elements, forming ternary to quinary phases. Binary phases were excluded because they have been extensively studied already, leaving little opportunity for novel materials discovery. For each structure–composition pair, we assigned the elements to specific prototype sites based on the given chemical formula. For example, randomly selecting the prototype “A2BC4_cF56_227_c_b_e-001” (a normal spinel in the AFLOW prototype library) and a composition of Mn–Fe–S would yield six spinel structures by exploring all possible arrangements of Mn, Fe, and S on the A/B/C sites. Charge balance is then assessed using common oxidation states provided by pymatgen.³⁷ If charge balance is plausible, the structure of the prototype is decorated accordingly, and the resulting materials undergo further evaluation using density functional theory (DFT). In the previous example, only Mn₂FeS₄ and Fe₂MnS₄ would be retained of the six enumerated structures. We note that while this method can produce new compositions, it will never produce an entirely new structure given its reliance on established prototypes. Our use of common oxidation states further limits the breadth of compositions that may be generated, though this requirement may be relaxed by the user if more exotic materials are desired.

Data-driven ion exchange

The Materials Project (MP) contains DFT-calculated properties for ~200 000 materials, providing a solid foundation for materials discovery campaigns.²⁹ Starting from these materials, we leveraged the data-mined substitution prediction (DMSP) algorithm³⁸ implemented in pymatgen³⁷ to replace one or more ions of a given compound to yield new hypothetical materials. The substitutions are guided by conditional probabilities, p_{DMSP} , which quantify the likelihood of substituting one ion for another while retaining the original crystal structure. These probabilities are derived from a probabilistic model trained on the inorganic crystal structure database (ICSD), an experimental database of known crystal structures.³⁹ Substitutions were performed for pairs of species with $p_{\text{DMSP}} > 0.001$ (the default value in pymatgen) to balance the tradeoff between generating novel substitutions and maintaining structural plausibility. Multiple substitutions were allowed per material, enabling both single and multi-site exchanges. For example, starting from CaTiO₃, yields substitutions like SrTiO₃ and SrZrO₃ (single-site) in addition to SrTiS₃ and SrZrS₃ (multi-site). DMSP was applied in two different modes: (1) to generate stable materials and (2) to generate materials with desired properties. For the first task, we randomly extracted stable parent structures from MP and substituted at least one ion with a species not already present in the original composition. To generate materials with desired properties, we selected materials having specified target values of that property (*e.g.*, band gaps near 3 eV) and performed ion substitution in a similar fashion to generate new materials. As with random enumeration, data-driven ion exchange can only produce

materials with novel compositions, whereas the structures will always be based on known materials.

Generative modeling

Four models were assessed in this work: CrystaLLM, FTCP, CDVAE, and MatterGen. CrystaLLM¹⁵ is a transformer-based large-language model designed to learn from tokenized representations of crystallographic information files (CIFs). We evaluated three pre-trained versions of this model (available at <https://github.com/lantunes/CrystaLLM>): two were trained on the MP-20 dataset,¹³ which contains 45 231 stable structures from the Materials Project,²⁹ using different model sizes (denoted small and large), and one was trained on a larger dataset of 2.3 million structures. These models require only a chemical formula as input, from which a CIF file is generated. All compositions fed to each instance of CrystaLLM were obtained through random enumeration of charge-balanced formulae. For consistent comparison with other models evaluated in this work, results shown in the main text (Fig. 1 and 2) correspond to the large model trained on MP-20; results from the other two versions are provided in the ESI.†

FTCP¹² encodes materials using real-space features (lattice vectors, one-hot encoded element vectors, site coordinates, and occupancies) and reciprocal-space features derived from a Fourier transform of elemental property vectors. Two FTCP-based autoencoders were trained on the MP-20 dataset: one conditioned on formation energy and electronic band gap and another conditioned on formation energy and bulk modulus. Due to the limited availability of elastic property data, the latter autoencoder was trained and validated on a subset of MP-20 containing 9361 materials. Materials were then generated by randomly sampling points in the latent space nearby known materials from the training set. Default hyperparameters supplied at <https://github.com/PV-Lab/FTCP> were used for the training and generation steps.

CDVAE¹³ combines a variational autoencoder (VAE) with a diffusion model to generate new materials. Sampling from the latent space predicts composition, lattice vectors, and the number of atoms in the unit cell, which are used to randomly initialize structures. The diffusion component of the model then “de-noises” these random structures by iteratively perturbing atoms toward equilibrium positions. We trained CDVAE on MP-20 without any conditioning of its latent space, allowing it to be used only for the generation of stable materials. During the generation step of each model, we did not place any constraints on the elements or symmetries that may be created. Default hyperparameters supplied at <https://github.com/txie-93/cdvae> were used for the training and generation steps.

MatterGen²² is a more recent diffusion model designed to create stable materials by jointly denoising atom types, coordinates, and lattice parameters through an equivariant score network. MatterGen can be fine-tuned for targeted properties but was used here in its base configuration (available at <https://github.com/microsoft/mattergen>) for unconstrained generation, re-trained on the MP-20 dataset for consistent comparison with the other models evaluated in this work. This model was

used to generate structures in batches of 128 using the provided command-line tool. A similar procedure was used to generate materials from a MatterGen model that was pre-trained on a much larger dataset (Alex-MP-20), with results provided in the ESI.†

Novelty assessment

A material was considered novel if it was not already present in MP (as of June 2025), thus ensuring it was excluded from the training data used for the generative models evaluated here. To determine this, we queried all entries in MP with the same composition as a proposed structure. For each resulting entry, structural similarity was assessed using the StructureMatcher tool in pymatgen,³⁷ which compares structures based on their lattice parameters, atomic positions, and symmetry. We used loose tolerance parameters for this comparison to account for slight variations in computed structures. This included a lattice parameter tolerance of 0.25 Å, a site tolerance defined as 40% of the average free length per atom, and an angular tolerance of 10°. If no matching composition or structure was identified in the MP database, the material was classified as novel.

Density functional theory calculations

Novel materials obtained from each generation approach described above were relaxed using the PBE exchange–correlation functional⁴⁰ within the projector augmented wave (PAW) method as implemented in the Vienna *Ab Initio* Simulation Package (VASP).^{41,42} These calculations used a plane-wave cut-off energy of 520 eV with a Γ -centered k -point grid spacing of 0.22 Å⁻¹. Convergence criteria were set to 10⁻⁶ eV for the electronic optimizations and 0.03 eV Å⁻¹ for the ionic relaxations. This relatively loose convergence criterion was determined to be sufficient for evaluating stability in high throughput, though a tighter criterion should be used if the objective is to precisely determine equilibrium atomic positions (ESI,† Fig. S1). Symmetry was turned off to ensure accurate treatment of distortions, and spin polarization was included for materials containing potentially magnetic elements (Ti, V, Cr, Mn, Fe, Co, Ni, Cu, Nb, Mo, Tc, Ru, Rh, Pd, Ag, Ta, W, Re, Os, Ir, Pt, Ce, Eu, Pr, Nd, Pm, Sm, Gd, Tb, Dy, Ho, Er, Tm, Yb, Np, and U). These elements were chosen to be consistent with the Materials Project. All magnetic moments were initialized in a ferromagnetic configuration for all such compounds. For materials containing 3d transition metals, Hubbard U corrections (+ U) were applied to account for on-site Coulomb interactions following the conventions used in the Materials Project: 4.0 eV for Mn, 3.9 eV for Fe, 3.7 eV for Co, 6.2 eV for Ni, and 5.3 eV for Cu.²⁹

Thermodynamic stability with respect to all known competing phases in MP was evaluated using the decomposition energy (ΔE_d).⁴³ For unstable materials with $\Delta E_d > 0$, this measure is equivalent to the energy above the convex hull (E_{hull}). It quantifies the energy difference between the proposed material and the lowest energy combination of competing phases. For stable materials with $\Delta E_d \leq 0$, the decomposition energy is the energy by which the proposed material lies below the existing convex hull (if the proposed material were not

included in its construction). Total energies acquired from DFT calculations were transformed into formation energies (ΔE_f) using the MaterialsProject2020 compatibility scheme, which accounts for GGA/GGA+ U mixing and implements elemental reference energy corrections as described in previous work.⁴⁴ Competing phases were identified by constructing a phase diagram for each chemical system using the PhaseDiagram module from pymatgen.³⁷ These phase diagrams included all entries from MP (as of June 2025) as well as the entries generated from each AI model or baseline approach, allowing for an evaluation of stability against both known and hypothetical phases.

Electronic band gaps were computed by analyzing the eigenvalue band properties obtained from VASP calculations using pymatgen,³⁷ with the band gap defined as the energy difference between the valence band maximum and conduction band minimum. The bulk modulus of each structure was computed by fitting a Birch–Murnaghan equation of state⁴⁵ to relaxed (but fixed volume) total energy calculations performed at seven volumes ranging from 97% to 103% of the equilibrium volume. These volumes were generated by isotropically scaling the lattice vectors of the relaxed equilibrium structures. The equilibrium bulk modulus and its pressure derivative were extracted from the fit, providing a measure of each material's resistance to volumetric deformation.

Machine learning filtering

Universal machine learning interatomic potentials (uMLIPs) offer an efficient way to screen large numbers of candidate materials obtained from any of the methods described above, enabling the rapid identification of promising structures prior to more computationally expensive DFT calculations.⁴⁶ Here, we used CHGNet³⁵ to compute the internal (0 K) energies of all candidate materials, which were then compared against the convex hulls in MP to assess thermodynamic stability. Each structure was relaxed using the atomic simulation environment

(ASE)⁴⁷ with CHGNet-based force fields, ensuring the forces acting on its atoms converged below $0.1 \text{ eV } \text{\AA}^{-1}$. This relatively loose convergence criterion was chosen to improve computational efficiency for the task of high-throughput screening. CHGNet was trained on GGA and GGA+ U calculations with MaterialsProject2020Compatibility corrections applied,⁴⁴ enabling direct comparison with MP energies. Thermodynamic stability was evaluated by constructing phase diagrams for the relevant chemical systems, incorporating all MP entries (as of June 2025) alongside generated candidate structures. Materials predicted to be stable within these phase diagrams were passed to DFT calculations for further validation.

For property-specific screening, we leveraged two CGCNN³⁶ models: one trained on 16 458 DFT-calculated band gaps and the other on 2041 bulk moduli from MP. When generating materials with a high bulk modulus, we applied an acceptance criterion of CGCNN-predicted bulk moduli exceeding 200 GPa. Analogously, when generating materials with a band gap near 3 eV, we selected candidates with CGCNN-predicted band gaps in the range of 2.8 to 3.2 eV.

Results

Stability and novelty

Each method described above was used to generate a sufficient number of candidates such that 500 novel materials – *i.e.*, ones that are not already present in the Materials Project (MP) – could be selected and passed to DFT calculations. The average time required by each method to generate these materials is listed in ESI,† Table S1. Their decomposition energies (ΔE_d), computed relative to the convex hull defined by MP entries, are shown in Fig. 1. Random enumeration of known structure prototypes with charge-balanced chemical formulae results in a wide range of decomposition energies with a median value (ΔE_d^{med}) of $409 \text{ meV atom}^{-1}$. Only 1.4% of the novel materials

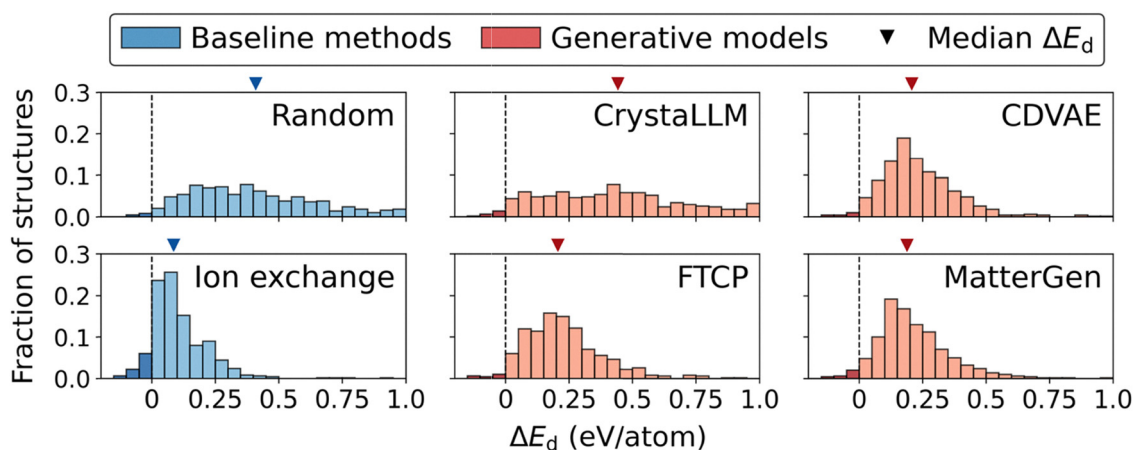


Fig. 1 Histograms showing DFT-computed decomposition energies (ΔE_d) of novel materials (not already present in the Materials Project) generated by two baseline methods and four generative models. For each of these six approaches, 500 materials were considered. The left column (blue) contains results from the baseline methods: random enumeration and ion exchange. The right columns (red) contain results from the generative models: CrystaLLM,¹⁵ FTCP,¹² CDVAE,¹³ and MatterGen.²² Triangular markers indicate median decomposition energies.

generated with this method are thermodynamically stable ($\Delta E_d \leq 0$). In contrast, a much higher stability rate can be achieved by leveraging analogies to known materials. Ion exchange results in a much tighter distribution of novel materials ($\Delta E_d^{\text{med}} = 85 \text{ meV atom}^{-1}$) close to the hull, with 9.2% of them being stable.

The high stability rate of ion exchange is impressive but perhaps unsurprising, given its proven efficacy in discovering new materials using high-throughput calculations performed over the past decade.^{48–53} This is especially true when leveraging known structure prototypes that can host a wide variety of compositions. For example, prior work has identified several hundred stable compositions in the perovskite structure through ion exchange of known materials.^{54–56} Similar approaches have also been applied to successfully uncover stable compositions in the spinel and delafossite structures,^{57,58} reinforcing the effectiveness of ion exchange for materials discovery.

Among the generative models, CrystaLLM produced novel materials with the widest range of energies ($\Delta E_d^{\text{med}} = 442 \text{ meV atom}^{-1}$) but a stability rate (2.4%) that is second only to MatterGen (3.0%). While both remain well below the stability rate of ion exchange (9.2%), the generative models provide unique flexibility in terms of capacity and training data. For example, the CrystaLLM model used in Fig. 1 had 200 million parameters and outperformed a smaller model variant with only 25 million parameters and a lower stability rate of 1.6% (ESI,† Fig. S2). Expanding the training set also improves performance, with CrystaLLM achieving a stability rate of 2.8% when trained on ~ 2.3 million structures (ESI,† Fig. S3). Similar improvements are possible with MatterGen, whose stability rate increases from 3.0% to 5.4% when trained on structures from the Alexandria database⁵⁹ in addition to those from MP-20 (ESI,† Fig. S4).

As both CDVAE and MatterGen are diffusion models, they result in a similar distribution of energies with ΔE_d^{med} of 207 and 188 meV atom^{-1} , respectively. These two models also produce comparable energy distributions to FTCP, a variational autoencoder, which yields $\Delta E_d^{\text{med}} = 205 \text{ meV atom}^{-1}$ and a moderate stability rate of 2.0%. Given the recent success of diffusion models, the competitiveness of FTCP (which is a VAE) is somewhat unexpected. We attribute this to its latent space sampling strategy, which biases generation toward materials that are structurally similar to known stable compounds. This approach enhances the model's stability rate, similar to the observed benefits of template-based strategies (such as ion exchange) in the baseline methods. However, it also comes at the cost of reduced novelty – 1309 materials were generated from FTCP before obtaining 500 materials not already present in MP (a novelty rate of $\sim 38.2\%$).

In this work, novelty is defined as a material being absent from MP. While this does not necessarily indicate the material has never been synthesized or is absent from all computational databases, it signifies that the material was not used for training of the generative models or as a template for ion exchange. All materials shown in Fig. 1 meet this definition

of novelty. However, the total number of generated materials required to obtain 500 novel ones varied across methods. The rate at which novel materials were produced by each method is listed in Table 1.

Between the two baseline methods, random enumeration yields a much higher novelty rate (98.6%) than ion exchange (72.4%). This reflects the unconstrained nature of random enumeration, which leads to the sampling of many previously unexplored chemical compositions. In contrast, our approach to ion exchange closely reflects traditional screening efforts,⁶⁰ and is therefore more likely to reproduce materials already present in computational databases such as MP. However, the use of ion exchange also comes with the benefit of generating more stable materials, resulting in a higher stability rate (9.2%) than random enumeration (1.4%).

Three of the generative models – CrystaLLM, CDVAE, and MatterGen – exhibit high novelty rates $>90\%$. In contrast, only 38.2% of the materials generated by FTCP are novel. This result is consistent with FTCP's strategy of sampling around known materials in its latent space, a factor that likely also contributes to its reasonably high stability rate. To assess the impact of the sampling strategy used by FTCP, we generated several new sets of materials at iteratively greater distances from known materials in its latent space. The results, shown in ESI,† Table S2 and Fig. S5, demonstrate that sampling further away from known materials leads to higher novelty rates (reaching 95%) but also lower stability rates ($\leq 1\%$). The clear inverse correlation between these two metrics underscores the tradeoff that exists between stability and novelty during materials discovery campaigns.

To more broadly assess novelty, we examined the fraction of generated materials absent from two additional sources: Alexandria,⁵⁹ a computational database containing over 4.5 million structures, and the ICSD,⁶¹ an experimental database with approximately 300 000 structures. As detailed in ESI,† Table S3, the novelty rate of each method decreased slightly upon comparing to these additional databases. However, the overall trends remain unchanged: random enumeration achieves the highest novelty rate (94.0%) while FTCP exhibits the lowest (35.0%). It is notable that even this lower novelty rate constitutes a substantial fraction of the generated materials, which suggests there remains ample opportunity for materials discovery even as these expansive databases continue to grow. While the novelty rate distribution is interesting, we argue that this is less important than the stability rates since novelty assessments are computationally inexpensive compared to stability assessments, which require DFT calculations.

It is worth noting that the stability rates of novel materials reported in this work are generally lower than the “SUN rates” (stability, uniqueness, and novelty) reported in prior work. For example, MatterGen and CDVAE have previously reported SUN rates of $\sim 38\%$ and $\sim 14\%$, respectively.²² Both values are much higher than the 1.8–3.0% stability rates found in our current study. This discrepancy arises primarily from differences in stability criteria. Previous work considered all materials within 100 meV atom^{-1} of the convex hull to be “stable.” Directly comparable metrics based on this definition are provided in ESI,†

Table 1 Stability and novelty rates of materials generated from each method. ΔE_d^{med} is the median decomposition energy of generated materials not already present in the Materials Project. Also listed is the prototype novelty rate, defined as the percentage of proposed materials whose structures cannot be indexed to a known prototype in the AFLOW database, and the stability rate of materials in these novel prototypes. The bold value in each column denotes the highest rate achieved among all methods. Statistics are based on 500 novel materials generated by each method

Method	ΔE_d^{med} (meV atom ⁻¹)	Stability rate (%)	Novelty rate (%)	Novel prototype rate (%)	Novel prototype stability rate (%)
Random	409	1.4	98.6	0	0
Ion exchange	85	9.2	72.4	0	0
CrystaLLM	442	2.4	98.2	1.0	0
CDVAE	207	1.8	96.0	8.2	0
FTCP	205	2.0	38.2	1.8	0
MatterGen	188	3.0	91.8	7.2	0

Table S4. However, we enforce a stricter definition of stability in the main text of our work, requiring that materials lie on the convex hull ($\Delta E_d \leq 0$) to be considered stable. While many previously synthesized materials are computed to be thermodynamically unstable with DFT, the probability a material can be synthesized is inversely proportional to the magnitude of this instability.^{62,63} It is difficult to define a general “rule-of-thumb” for accessible ΔE_d values, as this will depend on the nature of a material, its competing phases, and available synthetic routes.^{43,63,64} Materials computed to be on the hull are likely to be stable at ambient conditions, though synthesizing even hull-stable materials can be challenging.^{65–67} We argue that a stricter stability criterion is more meaningful, though a looser cutoff may still be appropriate if one is less concerned with the risk of false positives that would lead to unsuccessful synthesis attempts.

The generative models tested here may not lead to the highest stability rates, but they are unique in their ability to generate new structural frameworks that cannot be mapped to any known prototypes. This sets them apart from baseline methods, which rely entirely on existing templates and therefore exhibit 0% prototype novelty rates (Table 1). The generative models evaluated in this work achieve prototype novelty rates ranging from 1.0% (CrystaLLM) to 8.2% (CDVAE). As shown in ESI,† Fig. S6, a majority of these materials lie far above the convex hull, exhibiting $\Delta E_d > 100$ meV atom⁻¹. While those from MatterGen are closer to the convex hull on average, none achieve $\Delta E_d \leq 0$. The lack of proposed structures that are both stable and adopt novel prototypes highlights the need for generative models that can effectively balance thermodynamic stability with structural novelty.

Filtering stability with CHGNet

Filtering materials with uMLIPs such as CHGNet³⁵ provides a computationally efficient way to improve the stability rate of generation campaigns. Unlike DFT calculations, which are time-consuming and resource-intensive, CHGNet can be used to estimate the internal energy of a material within seconds. This energy can then be compared with a database of DFT-calculated energies to approximate thermodynamic stability. The efficiency of this method allows it to be integrated with any method for generating new materials, whether it be a baseline or generative model. In Fig. 2, we compare the cumulative distribution functions (CDFs) of DFT-calculated stability results without (left) and with CHGNet filtering (right). In the left

panel, materials obtained directly from each method were evaluated by computing their energies with DFT and comparing them to the MP convex hull. In the right panel, only materials predicted to be thermodynamically stable by CHGNet were included for subsequent DFT calculations.

The unfiltered results in the left panel of Fig. 2 serve as a reference to compare each method’s ability to generate stable materials. These are CDFs of the same histograms shown in Fig. 1. As observed in the prior section, ion exchange performs best in generating materials that are stable or close to the convex hull. This is evidenced by a steep rise in its CDF, positioned far to the left of all other methods. There is close competition among the next best three methods – MatterGen, CDVAE, and FTCP – whose CDFs overlap throughout a wide range of energies, reaching 80% near $\Delta E_d \approx 300$ meV atom⁻¹. In contrast, CrystaLLM and random enumeration yield CDFs that increase more gradually, reaching 80% only at energies above $\Delta E_d \approx 600$ meV atom⁻¹. This suggests that most materials produced by these two methods are unlikely to be accessed experimentally.^{63,64}

The right panel of Fig. 2 highlights the beneficial effect of CHGNet filtering, as the consistent leftward shift in all CDFs indicates a greater proportion of materials that are stable or close to the hull. However, results still vary substantially across different generation approaches. Filtered materials from ion exchange show a high stability rate of 15.2%, with the corresponding CDF reaching 80% at $\Delta E_d \approx 100$ meV atom⁻¹. Random enumeration also improves after filtering, achieving a stability rate of 7.6%. Among the generative models, CrystaLLM and FTCP benefit the most from CHGNet filtering, with stability rates increasing to 17.0% and 22.4%, respectively. In contrast, MatterGen and CDVAE show only modest gains, with relatively small shifts in their CDFs and updated stability rates ranging from 3.8% to 8.8%.

We speculate that filtering is less effective for the diffusion models as they often generate materials that fall outside of CHGNet’s training distribution – for example, in under-sampled chemistries or structures that are far out-of-equilibrium – potentially reducing the accuracy of stability predictions and limiting their performance gains. Indeed, ESI,† Fig. S8 shows large mean absolute errors (154 to 156 meV atom⁻¹) on structures from CDVAE and MatterGen. Large prediction errors (139 meV atom⁻¹) are also observed on structures generated through random enumeration. Despite adhering to known structure templates, random enumeration more often produces

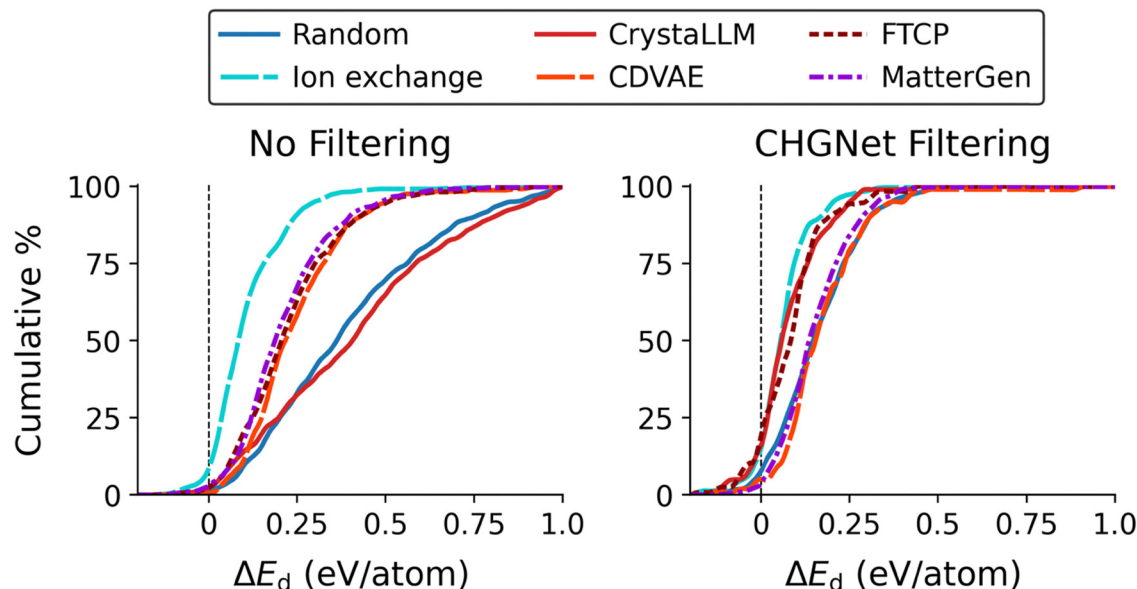


Fig. 2 Cumulative distribution functions (CDFs) showing the percentage of materials that satisfy a decomposition energy (ΔE_d) cutoff, with each line color-coded by the method used to generate these materials. The left panel displays CDFs for 500 novel materials generated directly by each method, including two baseline approaches (random enumeration and ion exchange) in blue tones and four generative models (CrystaLLM, CDVAE, FTCP, and MatterGen) in red/purple tones. The right panel displays CDFs for 500 novel materials filtered by CHGNet-predicted stability, including only those CHGNet predicts to have $\Delta E_d \leq 0$. Filtered energy distributions are also displayed in ESI,† Fig. S7.

exotic compositions with less representation in CHGNet's training set.

To assess the compositional diversity of generated materials, we provide heatmaps of element frequencies and histograms showing the number of elements per novel compound in ESI,† Fig. S9 and S10. Random enumeration produces compositions spanning much of the periodic table with relatively even distributions of pnictides, chalcogenides, and halides. However, these compositions are generally limited to ternary prototypes, reflecting the dominance of known three-element structures. In contrast, ion exchange produces a more diverse set of compounds – including quaternaries and quinarys – but the overall composition space is narrower, skewed toward oxides which are disproportionately prevalent in MP. This contributes to ion exchange providing the lowest CHGNet prediction error (47 meV atom⁻¹) among all methods.

Similarly, most of the generative models produce oxides at disproportionately high rates – reflecting bias in the MP-20 dataset on which they were all trained. For example, CDVAE generates a relatively narrow range of compositions, with 23.4% containing oxygen. However, it also produces more complex chemical formulae than template-based methods, with up to nine elements per compound. These multicomponent oxides are more often novel than compositions with fewer elements; however, they also compete with many compounds in the high-dimensional phase diagram, which contributes to the lower stability rate of CDVAE. Other methods generate a balance of binaries, ternaries, and quaternaries with broad periodic table coverage but a slight preference for oxides and halides. This further reflects bias in the MP-20 dataset used for training, which can be mitigated by expanding the set to include more

diverse compounds. For instance, the proportion of oxides generated by MatterGen drops from 21.2% to 10.9% after incorporating the Alexandria dataset into its training. FTCP demonstrates the most compositional diversity of the models tested here, sampling a wide range of elements with only 4.4% of its materials containing oxygen. Although FTCP is trained on the MP-20 dataset, we suspect its strategy of latent space sampling enables it to interpolate between known compounds and explore regions of composition space that are underrepresented in the training data.

Generating materials with targeted properties

Beyond generating thermodynamically stable materials, it is also useful to generate materials with targeted properties for particular applications. To this end, we adapted our baseline methods and re-trained a generative model to target desired band gaps and bulk moduli, aligning with the property-driven generation strategies demonstrated in recent work.²² First targeting materials with a band gap near 3 eV, two baseline methods were applied: random enumeration and ion exchange. Additionally, we applied ML filtering (based on predictions from CGCNN)³⁶ to the randomly enumerated materials to assess its impact on targeting specific properties. Finally, we tested one generative model, FTCP, whose latent space can be conditioned on specific properties. A total of 500 novel materials were generated from each method, and their distributions of computed band gaps are shown in Fig. 3.

Random enumeration produced a wide variety of materials, with 30.2% of the novel ones being metallic. Only 11.2% of these materials exhibited a band gap within 0.5 eV of the desired value (3 eV), demonstrating the low success rate of

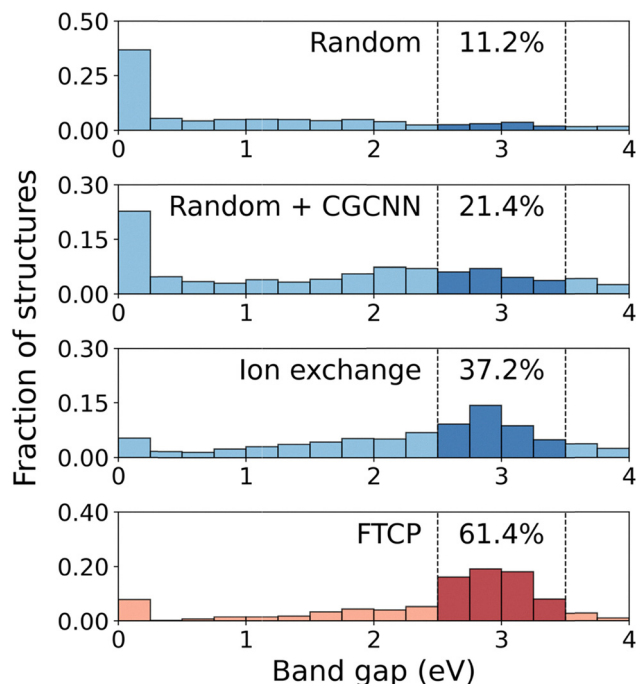


Fig. 3 Histograms showing density functional theory (DFT) computed band gap distributions of structures generated by two baseline methods (random enumeration and ion exchange, colored blue), CGCNN (also colored blue) applied to filter the randomly enumerated materials, and one generative model: FTCP (colored red). For each of the four approaches, 500 novel materials were considered. With the exception of random enumeration, all methods specifically targeted materials with a band gap near 3 eV. The percentage of generated materials with a band gap in the range of 2.5 to 3.5 eV is displayed above the shaded bars in each subplot.

computational screening when no guidance is provided. Applying CGCNN to filter these randomly enumerated materials improved the results considerably. By only retaining materials with CGCNN-predicted band gaps near 3 eV, the proportion of metals dropped to 18.8%, and 21.4% of the filtered materials exhibited band gaps within 0.5 eV of the target. As with CHGNet-filtering, this showcases the utility of ML-based screening for quickly refining large pools of candidate materials.

Data-driven ion exchange performed even better than CGCNN filtering of randomly enumerated materials, leveraging its ability to generate hypothetical compounds by substituting ions in known materials from MP that already have band gaps close to 3 eV. This method resulted in only 5.6% of the novel materials being metallic and a substantial 37.2% of them having a band gap within 0.5 eV of the target. This strong performance may not be entirely surprising as many of the compositional changes introduced by ion exchange are relatively minor, especially when the substituted element constitutes a small fraction of the overall chemical formula. This mirrors our findings from the previous section, highlighting the tradeoff between achieving success – whether in targeted properties or stability – and prioritizing novelty or diversity in the generated structures.

FTCP outperformed all other methods in targeting electronic band gap, with 61.4% of its novel materials exhibiting a

band gap within 0.5 eV of the desired value (3 eV). This success likely stems from FTCP's latent space sampling informed by known compounds with band gaps close to the target, which enables the generation of materials with structural or compositional similarities to the reference points. Thermodynamic stability remains an important consideration, as only 3.0% of the novel materials generated by FTCP are stable, compared with 15.2% of those generated by ion exchange (ΔE_d distributions provided in ESI,† Fig. S11).

Using the same four methods described above (for targeting a desired band gap), we next generated materials with the objective of maximizing bulk modulus. This task fundamentally differs from the previous band gap-related objective by focusing on materials with extreme properties (*e.g.*, maximal bulk modulus) instead of those within an intermediate range (*e.g.*, band gaps near 3 eV). A total of 500 materials were sampled from each method, and their bulk moduli were computed using Birch–Murnaghan equations of state fit to DFT-computed energies. The resulting distributions of bulk moduli are shown in Fig. 4. Materials generated through random enumeration follow a Poisson-like distribution of bulk moduli with a peak near 50–60 GPa, closely resembling the known distribution of elastic properties for materials in MP.⁶⁸ If we define success as finding novel materials with a bulk modulus ≥ 300 GPa, then random enumeration achieves this at a rate of only 3.0%. When CGCNN

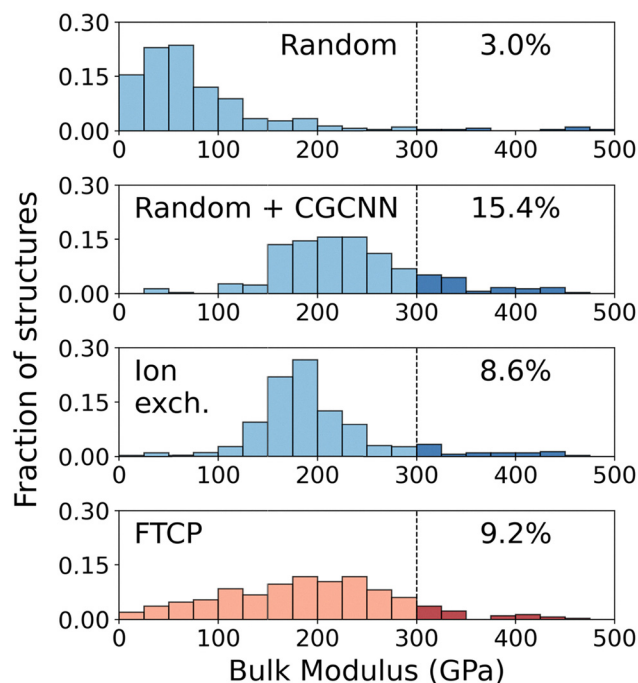


Fig. 4 Histograms showing density functional theory (DFT) computed bulk moduli distributions of structures generated by two baseline methods (random enumeration and ion exchange, colored blue), CGCNN (also colored blue) applied to filter the randomly enumerated materials, and one generative model: FTCP (colored red). For each of the four approaches, 500 novel materials were considered. With the exception of random enumeration, all methods specifically targeted materials with a bulk modulus ≥ 300 GPa. The percentage of generated materials satisfying this objective is displayed above the shaded bars in each subplot.

is applied to filter these materials, it causes a noticeable shift in the distribution toward higher bulk moduli, and 15.4% of the filtered materials exhibit a bulk modulus ≥ 300 GPa.

When applied to known materials in MP with high bulk moduli, ion exchange performs more modestly, with 8.6% of the resulting materials exhibiting a bulk modulus ≥ 300 GPa. This smaller shift in the distribution likely reflects the tendency for ion exchange to introduce only minor compositional changes, which limits its ability to substantially alter the mechanical properties of the original materials – many of which (in MP) do not exhibit anomalously high bulk moduli. FTCP performed slightly better than ion exchange but worse than CGCNN-based filtering of randomly enumerated materials, with 9.2% of the compounds generated by FTCP exhibiting a bulk modulus ≥ 300 GPa.

Compared to its strong performance on electronic band gap, we suspect FTCP is less effective here given the scarcity of materials with extremely high bulk moduli in MP. This lack of training data may limit the conditioning of the autoencoder's latent space on extreme bulk modulus values. FTCP also yields a low stability rate of 2.0% when targeting novel materials with high bulk modulus. CGCNN filtering and ion exchange face similar limitations, with stability rates of 2.0% and 1.8%, respectively (ΔE_d distributions provided in ESI,† Fig. S12). These uniformly low percentages across all evaluated methods highlight the challenge of identifying “exceptional” materials, as the inherent scarcity of analogs in the materials space and limited training data inhibit the development of effective models for both generation and filtering.⁶⁹

Discussion

The recent surge in generative models for inorganic crystalline materials underscores the growing need for benchmarks to assess their performance. While metrics such as stability and novelty provide valuable insight, there is a lack of clear baselines for comparison. To address this gap, we developed and evaluated two baseline methods: random enumeration of charge-balanced compositions in known structure prototypes, and targeted ion substitution of known materials with desired properties. These approaches leverage existing data from AFLOW^{33,34} and the Materials Project,²⁹ which offer a wealth of information on structure prototypes and calculated properties acquired from DFT calculations. They also benefit from simple yet powerful chemical heuristics; charge balance favors validity of proposed chemical formulae, and substitutions are performed on ions of comparable size and oxidation state. As a result, the baseline methods perform well in generating stable materials not found in existing databases. Random enumeration achieves a modest stability rate of 1.4%, while ion exchange achieves a much higher rate of 9.2%. There remains ample opportunity to further increase these rates as additional chemical heuristics are introduced to better filter computationally proposed materials.⁷⁰

The strong performance of the baseline methods establishes a high benchmark for generative models to meet or exceed. For

this task, we tested a variational autoencoder (FTCP),¹² a large-language model (CrystaLLM),¹⁵ and two diffusion models (CDVAE and MatterGen).^{13,22} Our tests showed MatterGen to be most effective in generating materials on or close to the hull, though its stability rate of 3.0% still falls well below that of ion exchange (9.2%). Nevertheless, generative models excel in generating materials with a high degree of structural novelty; up to 8.2% cannot be mapped to any known structure prototype in the AFLOW database. The capability of generating entirely new structural arrangements is unique to the generative models, but their low stability rates leave much room for improvement. One promising direction is to expand the training data for these models. For example, we found that MatterGen achieves a higher stability rate of 5.4% when trained on materials from Alexandria in addition to MP-20. Similar improvements were observed for CrystaLLM.

It is important to note that many of the comparisons made in this work depend on the stability threshold used to define success. We adopted a strict criterion requiring materials to lie on the convex hull ($\Delta E_d \leq 0$) to be considered stable, which reduces false positives but also penalizes the generation of near-stable candidates that may be synthesizable. A looser threshold would raise the stability rates across all methods and shift the relative performance of each approach. For example, a threshold of $\Delta E_d \leq 100$ meV atom⁻¹ increases the stability rate of the generative models to 11–20%, but these rates still far well below ion exchange, which achieves a stability rate of 58% using the same threshold (ESI,† Table S4).

In addition to generating a large proportion of materials near the convex hull, generative models also perform well in targeting specific properties when sufficient training data is available. For instance, FTCP achieves a high success rate of 61.4% in generating materials with a desired band gap near 3 eV, far surpassing the 37.2% achieved by ion exchange. This performance diminishes when targeting extreme values of properties such as high bulk moduli (> 300 GPa) that are less well represented in the training set. However, improved results can likely be obtained by running additional calculations on materials with extreme property and feeding them back into the generative models as training data.

To enhance the performance of the methods discussed in this paper, machine learning models were used to filter proposed materials based on predicted thermodynamic stability or desired properties. Our results demonstrate that this is a highly effective approach. Filtering by predicted stability using a pre-trained uMLIP (CHGNet)³⁵ substantially improves the stability rates of generated materials. A notably high 22.4% of novel materials generated by FTCP lie on the DFT convex hull after filtering. This performance boost achieved by filtering is diminished for some generative models like CDVAE and MatterGen, which produce more exotic materials that fall outside of CHGNet's training distribution. However, uMLIPs are likely to become more effective at filtering such materials as the breadth and diversity of their training data improves.⁷¹ This trend is evident in the correlation between prediction error and the filtered stability rate (ESI,† Fig. S13), suggesting that reducing

prediction error in future uMLIPs should further improve stability rates.

Similar findings were observed when using a pre-trained graph neural network (CGCNN)³⁶ to filter materials by predicted band gap and bulk modulus. Doing so leads to a near three-fold increase in the success rate of identifying materials with desired properties compared to random enumeration but remains relatively low (15.4%) when targeting extreme property values (e.g., a bulk modulus >300 GPa). It also leads to a decrease in the stability rate of the proposed materials, though incorporating a uMLIP-based stability filter could mitigate this issue. As with uMLIPs, these findings underscore the need to broaden and diversify training data for property prediction models to enhance the efficiency of generative approaches in identifying novel materials with exceptional properties.

Our findings demonstrate that there is still room for improvement in the design of generative models for inorganic materials, particularly when they are used to find new materials that are thermodynamically stable. To streamline the development of future models, we provide all of the data and code from this work in a publicly accessible GitHub repository (see Data availability statement). We envision these resources being used for benchmarking generative models and integrating them with traditional screening methods to enhance the success rate in discovering new materials that are likely to be synthesized and display desired properties.

Conflicts of interest

There are no conflicts to declare.

Data availability

The code for generating materials through random enumeration and ion exchange, in addition to filtering these materials on the basis of machine learning predictions, is available at https://github.com/Bartel-Group/matgen_baselines. This repository also includes CIF files for the structures generated by each method, as well as pre-trained models on the MP-20 dataset and installation instructions for CrystaLLM, FTCP, CDVAE, and MatterGen.

Acknowledgements

This work was supported by the 3DEAP NRT, NSF grant no. 2345719, and new faculty start-up funds from the University of Minnesota. The authors also acknowledge the Minnesota Supercomputing Institute (MSI) at the University of Minnesota for providing resources that contributed to the research results reported herein. This work was also enabled by the dedication and open-source contributions of prior developers, including those behind the pre-trained graph neural networks (CHGNet and CGCNN), generative models (CrystaLLM, CDVAE, FTCP, and MatterGen), structure prediction algorithms (pymatgen),

structure prototypes (AFLOW), and DFT databases (Materials Project).

References

- 1 A. K. Cheatham, R. Seshadri and F. Wudl, Chemical synthesis and materials discovery, *Nat. Synth.*, 2022, **1**, 514–520.
- 2 A. Manthiram, J. C. Knight, S.-T. Myung, S.-M. Oh and Y.-K. Sun, Nickel-Rich and Lithium-Rich Layered Oxide Cathodes: Progress and Perspectives, *Adv. Energy Mater.*, 2016, **6**, 1501010.
- 3 K. Nomura, *et al.*, Thin-Film Transistor Fabricated in Single-Crystalline Transparent Oxide Semiconductor, *Science*, 2003, **300**, 1269–1272.
- 4 K. M. Shen and J. C. Seamus Davis, Cuprate high-Tc superconductors, *Mater. Today*, 2008, **11**, 14–21.
- 5 H. Park, Z. Li and A. Walsh, Has generative artificial intelligence solved inverse materials design?, *Matter*, 2024, **7**, 2355–2367.
- 6 Z. Wang, H. Hua, W. Lin, M. Yang and K. C. Tan, *Crystalline Material Discovery in the Era of Artificial Intelligence*, *arXiv*, 2025, preprint, arXiv:2408.08044, DOI: [10.48550/arXiv.2408.08044](https://doi.org/10.48550/arXiv.2408.08044).
- 7 A. Noura, N. Sokolovska and J.-C. Crivello, *CrystalGAN: Learning to Discover Crystallographic Structures with Generative Adversarial Networks*, *arXiv*, 2018, preprint, arXiv:1810.11203, DOI: [10.48550/1810.11203](https://doi.org/10.48550/1810.11203).
- 8 S. Kim, J. Noh, G. Gu, A. Aspuru-Guzik and Y. Jung, Generative Adversarial Networks for Crystal Structure Prediction, *ACS Cent. Sci.*, 2020, **6**, 1412–1420.
- 9 Y. Zhao, *et al.*, High-Throughput Discovery of Novel Cubic Crystal Materials Using Deep Generative Neural Networks, *Adv. Sci.*, 2021, **8**, 2100566.
- 10 J. Noh, *et al.*, Inverse Design of Solid-State Materials *via* a Continuous Representation, *Matter*, 2019, **1**, 1370–1384.
- 11 C. J. Court, B. Yildirim, A. Jain and J. M. Cole, 3-D Inorganic Crystal Structure Generation and Property Prediction *via* Representation Learning, *J. Chem. Inf. Model.*, 2020, **60**, 4518–4535.
- 12 Z. Ren, *et al.*, An invertible crystallographic representation for general inverse design of inorganic crystals with targeted properties, *Matter*, 2022, **5**, 314–335.
- 13 T. Xie, X. Fu, O.-E. Ganea, R. Barzilay and T. Jaakkola, *Crystal Diffusion Variational Autoencoder for Periodic Material Generation*, *arXiv*, 2022, preprint, arXiv:2110.06197, DOI: [10.48550/arXiv.2110.06197](https://doi.org/10.48550/arXiv.2110.06197).
- 14 R. Zhu, W. Nong, S. Yamazaki and K. Hippalgaonkar, WyCryst: Wyckoff inorganic crystal generator framework, *Matter*, 2024, **7**, 3469–3488.
- 15 L. M. Antunes, K. T. Butler and R. Grau-Crespo, Crystal structure generation with autoregressive large language modeling, *Nat. Commun.*, 2024, **15**, 10570.
- 16 D. Flam-Shepherd and A. Aspuru-Guzik, *Language models can generate molecules, materials, and protein binding sites directly in three dimensions as XYZ, CIF, and PDB files*, *arXiv*, 2023, preprint, arXiv:2305.05708, DOI: [10.48550/arXiv.2305.05708](https://doi.org/10.48550/arXiv.2305.05708).

- 17 N. Gruver, *et al.*, *Fine-Tuned Language Models Generate Stable Inorganic Materials as Text*, *arXiv*, 2024, preprint, arXiv:2402.04379, DOI: [10.48550/arXiv.2402.04379](https://doi.org/10.48550/arXiv.2402.04379).
- 18 P.-P. D. Breuck, H. A. Piracha, G.-M. Rignanese and M. A. L. Marques, *A generative material transformer using Wyckoff representation*, *arXiv*, 2025, preprint, arXiv:2501.16051, DOI: [10.48550/arXiv.2501.16051](https://doi.org/10.48550/arXiv.2501.16051).
- 19 T. Mohanty, M. Mehta, H. M. Sayeed, V. Srikumar and T. D. Sparks, *CrysText: A Generative AI Approach for Text-Conditioned Crystal Structure Generation using LLM*, *Chemrxiv*, 2024, preprint, DOI: [10.26434/chemrxiv-2024-gjhpq](https://doi.org/10.26434/chemrxiv-2024-gjhpq).
- 20 M. Alverson, *et al.*, *Generative adversarial networks and diffusion models in material discovery*, *Digital Discovery*, 2024, **3**, 62–80.
- 21 R. Jiao, *et al.*, *Crystal Structure Prediction by Joint Equivariant Diffusion*, *arXiv*, 2024, preprint, arXiv:2309.04475, DOI: [10.48550/arXiv.2309.04475](https://doi.org/10.48550/arXiv.2309.04475).
- 22 C. Zeni, *et al.*, *A generative model for inorganic materials design*, *Nature*, 2025, **639**, 624–632.
- 23 S. Yang, *et al.*, *Scalable Diffusion for Materials Generation*, *arXiv*, 2024, preprint, arXiv:2311.09235, DOI: [10.48550/arXiv.2311.09235](https://doi.org/10.48550/arXiv.2311.09235).
- 24 P. Hoellmer, *et al.*, *Open Materials Generation with Stochastic Interpolants*, *arXiv*, 2025, preprint, arXiv:2502.02582, DOI: [10.48550/arXiv.2502.02582](https://doi.org/10.48550/arXiv.2502.02582).
- 25 B. K. Miller, R. T. Q. Chen, A. Sriram and B. M. Wood, *FlowMM: Generating Materials with Riemannian Flow Matching*, *arXiv*, 2024, preprint, arXiv:2406.04713, DOI: [10.48550/arXiv.2406.04713](https://doi.org/10.48550/arXiv.2406.04713).
- 26 A. Sriram, B. K. Miller, R. T. Q. Chen and B. M. Wood, *FlowLLM: Flow Matching for Material Generation with Large Language Models as Base Distributions*, *arXiv*, 2024, preprint, arXiv:2410.23405, DOI: [10.48550/arXiv.2410.23405](https://doi.org/10.48550/arXiv.2410.23405).
- 27 L. Ruple, L. Torresi, H. Schopmans and P. Friederich, *Symmetry-Aware Bayesian Flow Networks for Crystal Generation*, *arXiv*, 2025, preprint, arXiv:2502.03146, DOI: [10.48550/arXiv.2502.03146](https://doi.org/10.48550/arXiv.2502.03146).
- 28 K. Tangsongcharoen, *et al.*, *Ref. 28: the article number has been derived from the provided DOI. Please check and correct as necessary via*, *arXiv*, 2025, preprint, arXiv:2501.08998, DOI: [10.48550/arXiv.2501.08998](https://doi.org/10.48550/arXiv.2501.08998).
- 29 A. Jain, *et al.*, *Commentary: The Materials Project: A materials genome approach to accelerating materials innovation*, *APL Mater.*, 2013, **1**, 011002.
- 30 H. Park, A. Onwuli and A. Walsh, *Exploration of crystal chemical space using text-guided generative artificial intelligence*, *Nat. Commun.*, 2025, **16**, 4379.
- 31 S. G. Baird, H. M. Sayeed, J. Montoya and T. D. Sparks, *matbench-genmetrics: A Python library for benchmarking crystal structure generative models using time-based splits of Materials Project structures*, *J. Open Source Softw.*, 2024, **9**, 5618.
- 32 J. Riebesell, *et al.*, *A framework to evaluate machine learning crystal stability predictions*, *Nat. Mach. Intell.*, 2025, **7**, 836–847.
- 33 M. J. Mehl, *et al.*, *The AFLOW Library of Crystallographic Prototypes: Part 1*, *Comput. Mater. Sci.*, 2017, **136**, S1–S828.
- 34 D. Hicks, *et al.*, *The AFLOW Library of Crystallographic Prototypes: Part 2*, *Comput. Mater. Sci.*, 2019, **161**, S1–S1011.
- 35 B. Deng, *et al.*, *CHGNet as a pretrained universal neural network potential for charge-informed atomistic modelling*, *Nat. Mach. Intell.*, 2023, **5**, 1031–1041.
- 36 T. Xie and J. C. Grossman, *Crystal Graph Convolutional Neural Networks for an Accurate and Interpretable Prediction of Material Properties*, *Phys. Rev. Lett.*, 2018, **120**, 145301.
- 37 S. P. Ong, *et al.*, *Python Materials Genomics (pymatgen): A robust, open-source python library for materials analysis*, *Comput. Mater. Sci.*, 2013, **68**, 314–319.
- 38 G. Hautier, C. Fischer, V. Ehrlicher, A. Jain and G. Ceder, *Data Mined Ionic Substitutions for the Discovery of New Compounds*, *Inorg. Chem.*, 2011, **50**, 656–663.
- 39 I. Levin, *NIST Inorganic Crystal Structure Database (ICSD)*. National Institute of Standards and Technology 10.18434/M32147 (2020).
- 40 J. P. Perdew, K. Burke and M. Ernzerhof, *Generalized Gradient Approximation Made Simple*, *Phys. Rev. Lett.*, 1996, **77**, 3865–3868.
- 41 G. Kresse and J. Furthmüller, *Efficiency of ab-initio total energy calculations for metals and semiconductors using a plane-wave basis set*, *Comput. Mater. Sci.*, 1996, **6**, 15–50.
- 42 G. Kresse and J. Furthmüller, *Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set*, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 1996, **54**, 11169–11186.
- 43 C. J. Bartel, *Review of computational approaches to predict the thermodynamic stability of inorganic solids*, *J. Mater. Sci.*, 2022, **57**, 10475–10498.
- 44 A. Wang, *et al.*, *A framework for quantifying uncertainty in DFT energy corrections*, *Sci. Rep.*, 2021, **11**, 15496.
- 45 F. Birch, *Finite Elastic Strain of Cubic Crystals*, *Phys. Rev.*, 1947, **71**, 809–824.
- 46 R. Jacobs, *et al.*, *A practical guide to machine learning interatomic potentials – Status and future*, *Curr. Opin. Solid State Mater. Sci.*, 2025, **35**, 101214.
- 47 A. H. Larsen, *et al.*, *The atomic simulation environment—a Python library for working with atoms*, *J. Phys.: Condens. Matter*, 2017, **29**, 273002.
- 48 X. Zhang, V. Stevanović, M. d’Avezac, S. Lany and A. Zunger, *Prediction of A₂BX₄ metal–chalcogenide compounds via first-principles thermodynamics*, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 2012, **86**, 014109.
- 49 D. W. Davies, *et al.*, *Computational Screening of All Stoichiometric Inorganic Materials*, *Chem*, 2016, **1**, 617–627.
- 50 P. Gorai, A. Ganose, A. Faghaninia, A. Jain and V. Stevanović, *Computational discovery of promising new n-type dopable ABX Zintl thermoelectric materials*, *Mater. Horiz.*, 2020, **7**, 1809–1818.
- 51 H.-C. Wang, S. Botti and M. A. L. Marques, *Predicting stable crystalline compounds using chemical similarity*, *npj Comput. Mater.*, 2021, **7**, 12.
- 52 A. Vasylenko, *et al.*, *Element selection for crystalline inorganic solid discovery guided by unsupervised machine learning of experimentally explored chemistry*, *Nat. Commun.*, 2021, **12**, 5561.

- 53 A. Merchant, *et al.*, Scaling deep learning for materials discovery, *Nature*, 2023, **624**, 80–85.
- 54 A. A. Emery, J. E. Saal, S. Kirklin, V. I. Hegde and C. Wolverton, High-Throughput Computational Screening of Perovskites for Thermochemical Water Splitting Applications, *Chem. Mater.*, 2016, **28**, 5621–5634.
- 55 Y. Wang, B. Baldassarri, J. Shen, J. He and C. Wolverton, Landscape of Thermodynamic Stabilities of A₂BB'O₆ Compounds, *Chem. Mater.*, 2024, **36**, 6816–6830.
- 56 J. Schmidt, *et al.*, Predicting the Thermodynamic Stability of Solids Combining Density Functional Theory and Machine Learning, *Chem. Mater.*, 2017, **29**, 5090–5103.
- 57 V. Kocovski, G. Pilania and B. P. Uberuaga, High-throughput investigation of the formation of double spinels, *J. Mater. Chem. A*, 2020, **8**, 25756–25767.
- 58 J. Shi, *et al.*, High-throughput search of ternary chalcogenides for p-type transparent electrodes, *Sci. Rep.*, 2017, **7**, 43179.
- 59 J. Schmidt, *et al.*, Improving machine-learning models in materials science through large datasets, *Mater. Today Phys.*, 2024, **48**, 101560.
- 60 K. T. Butler, J. M. Frost, J. M. Skelton, K. L. Svane and A. Walsh, Computational materials design of crystalline solids, *Chem. Soc. Rev.*, 2016, **45**, 6138–6146.
- 61 D. Zagorac, H. Müller, S. Ruehl, J. Zagorac and S. Rehme, Recent developments in the Inorganic Crystal Structure Database: theoretical crystal structure data and related features, *J. Appl. Crystallogr.*, 2019, **52**, 918–925.
- 62 A. Narayan, *et al.*, Computational and experimental investigation for new transition metal selenides and sulfides: The importance of experimental verification for stability, *Phys. Rev. B*, 2016, **94**, 045105.
- 63 W. Sun, *et al.*, The thermodynamic scale of inorganic crystalline metastability, *Sci. Adv.*, 2016, **2**, e160022.
- 64 M. Aykol, S. S. Dwaraknath, W. Sun and K. A. Persson, Thermodynamic limit for synthesis of metastable inorganic materials, *Sci. Adv.*, 2018, **4**, eaaq014.
- 65 A. Wustrow, *et al.*, Synthesis and Characterization of MgCr₂S₄ Thiospinel as a Potential Magnesium Cathode, *Inorg. Chem.*, 2018, **57**, 8634–8638.
- 66 A. Miura, *et al.*, Selective metathesis synthesis of MgCr₂S₄ by control of thermodynamic driving forces, *Mater. Horiz.*, 2020, **7**, 1310–1316.
- 67 S. Pandey, *et al.*, Steam-Assisted Ammonolysis of MoO₂ as a Synthetic Pathway to Oxygenated δ-MoN, *Materials*, 2025, **18**, 2340.
- 68 M. de Jong, *et al.*, Charting the complete elastic properties of inorganic crystalline compounds, *Sci. Data*, 2015, **2**, 150009.
- 69 J. Schrier, A. J. Norquist, T. Buonassisi and J. Brgoch, In Pursuit of the Exceptional: Research Directions for Machine Learning in Chemical and Materials Science, *J. Am. Chem. Soc.*, 2023, **145**, 21699–21716.
- 70 B. Das, K. Ji, F. Sheng, K. M. McCall and T. Buonassisi, Embedding human knowledge in material screening pipeline as filters to identify novel synthesizable inorganic materials, *Faraday Discuss.*, 2025, **256**, 587–600.
- 71 L. Barroso-Luque, *et al.*, *Open Materials 2024 (OMat24) Inorganic Materials Dataset and Models*, *arXiv*, 2024, preprint, arXiv:2410.12771, DOI: [10.48550/arXiv.2410.12771](https://doi.org/10.48550/arXiv.2410.12771).