
Space Group Equivariant Crystal Diffusion

**Rees Chang^{1*}, Angela Pak¹, Alex Guerra², Ni Zhan², Nick Richardson²,
Elif Ertekin³, Ryan P. Adams²**

¹Department of Materials Science & Engineering, University of Illinois at Urbana-Champaign

²Department of Computer Science, Princeton University

³Materials Research Laboratory, University of Illinois at Urbana-Champaign

Abstract

Accelerating inverse design of crystalline materials with generative models has significant implications for a range of technologies. Unlike other atomic systems, 3D crystals are invariant to discrete groups of isometries called the space groups. Crucially, these space group symmetries are known to heavily influence materials properties. We propose SGEquiDiff, a crystal generative model which naturally handles space group constraints with space group invariant likelihoods. SGEquiDiff consists of an $SE(3)$ -invariant, telescoping discrete sampler of crystal lattices; permutation-invariant, transformer-based autoregressive sampling of Wyckoff positions, elements, and numbers of symmetrically unique atoms; and space group equivariant diffusion of atomic coordinates. We show that space group equivariant vector fields automatically live in the tangent spaces of the Wyckoff positions. SGEquiDiff achieves state-of-the-art performance on standard benchmark datasets as assessed by quantitative proxy metrics and quantum mechanical calculations. Our code is available at <https://github.com/rees-c/sgequidiff>.

1 Introduction

Crystals comprise critical technologies like batteries [66], topological materials [90], electronic devices [96], photovoltaics [28], and more. Materials scientists have catalogued $\mathcal{O}(10^5)$ crystals experimentally [7] and $\mathcal{O}(10^6)$ *in silico* with density functional theory (DFT) simulation [15, 39, 79]. In contrast, the number of stable crystalline materials with five elements or less is estimated to exceed 10^{13} , and even higher-order compositions are common in real materials [16]. Generative models offer a promising path to rapidly explore the vast space of crystals [100, 40, 62, 10].

Unlike molecules, crystals span the periodic table and exhibit discrete spatial symmetries according to one of 230 *space groups* [2]. Specifically, crystals have invariances to discrete translations, rotations, reflections, and sequences thereof that transform atoms into themselves or into identical atoms. The list of space group actions which map a point into itself is called a *stabilizer group*. When a set of points in \mathbb{R}^3 have conjugate stabilizer groups, the set is called a *Wyckoff position*. Importantly, as shown in Figure 1, Wyckoff positions can have zero volume, comprising points, lines, or planes. These zero volume sets are referred to as *special Wyckoff positions*. We give a more formal treatment of space groups and Wyckoff positions in Sec. 3.

Despite the fact that most existing crystal generative models ignore space groups and Wyckoff positions, they are critical for modeling real materials. Firstly, space groups and Wyckoff positions correlate strongly with materials properties; Neumann’s principle states that all crystal properties share the same invariances as the crystal itself [65]. Thus even slightly perturbing atoms out of special Wyckoff positions will reduce the crystal’s space group symmetry and can subsequently cause significant (even discontinuous) changes to its macroscopic properties [14, 11, 82, 9]. Secondly, we

*Correspondence to reeswc2@illinois.edu

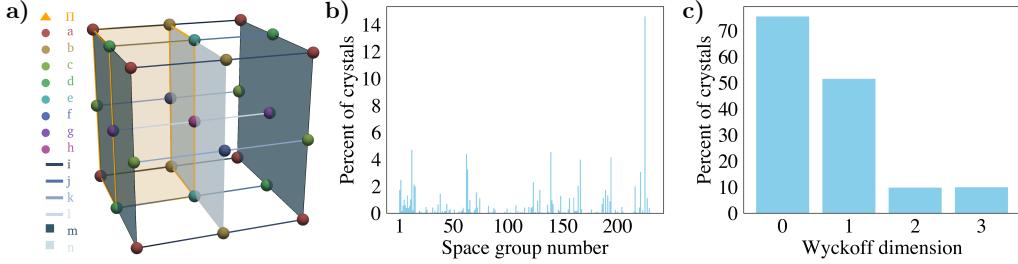


Figure 1: (a) The asymmetric unit (II) and special Wyckoff positions labeled by letter in the conventional unit cell of space group 10. (b-c) Histograms of occupied space groups and Wyckoff dimensionalities by crystals in the MP20 [100, 39, 7] training dataset. Space groups and Wyckoff positions were determined by the SpaceGroupAnalyzer module in pymatgen [70, 91] using tolerances of 0.1 \AA and 5° . These tolerances help account for the moderate convergence criteria of the Materials Project DFT relaxations.

show empirically in Figure 1 that known materials are usually invariant to high symmetry space groups with atoms in zero-volume Wyckoff positions. An intuitive explanation for this phenomenon is that atoms of the same species in a given crystal are most stable in a specific neighboring environment, and the tendency to attain this environment naturally leads to symmetry [32]. Yet, most existing crystal generative models learn continuous distributions over all three spatial dimensions of atom positions, assigning zero probability measure to placing atoms in special Wyckoff positions. One might argue that crystals from these models can be relaxed into high symmetry positions with DFT or machine learned force fields. However, besides the apparent difficulty of generating atoms sufficiently close to high symmetry positions for relaxation [107, 31, 42], such a framework assigns different model probabilities to crystals that relax into the same structure, obfuscating training and evaluation.

In this paper, we propose SGEquiDiff, which enforces space group constraints during generation and yields explicit space group invariant model likelihoods. SGEquiDiff combines SE(3)-invariant, discretized lattice sampling; permutation-invariant, transformer-based generation of occupied Wyckoff positions and elements; and space group equivariant diffusion of atomic coordinates. For the latter, we built a space group equivariant graph neural network and trained it with space group equivariant scores from a *Space Group Wrapped Normal* distribution.

The main contributions of our work are summarized as follows:

- We built a space group equivariant denoising graph neural network and regressed it against scores from a novel *Space Group Wrapped Normal* distribution.
- We prove that space group equivariant vector fields live in the tangent space of each Wyckoff position, ensuring atoms never leave their Wyckoff positions during diffusion. This result obviates the need to project atoms onto Wyckoff positions, which can lead to indeterminate probability distributions [41] or a discontinuous generation process [55, 78].
- SGEquiDiff proposes an explicit-likelihood autoregressive alternative to diffusion-based sampling of space group-constrained lattice parameters. The autoregressive sampling also enables masked in-filling. This may be useful, e.g., to generate optimal substrate materials with low lattice mismatch for epitaxial crystal growth [110, 19], commonly used for semiconductors and magnetic devices.
- We verified the efficacy of our model on standard crystal datasets with proxy metrics and rigorous DFT calculations.

2 Related Work

Early crystal generative models represented crystals as voxelized images [67, 34] or padded tensors of 3D coordinates [49, 75] to train variational autoencoders (VAE) [18] or generative adversarial networks [26]. Several works have enforced the SE(3) and periodic translational invariance of crystals by leveraging graph neural networks. One popular approach is to use diffusion models [83, 84, 33]

on crystal lattices, atom types, and atom positions [100, 40, 107, 57]. These models have also been extended with the flow matching framework [58], accelerating sampling [62]; the stochastic interpolants framework [1], expanding the model design space [35]; and the Bayesian Flow Networks framework [27, 97] which iteratively refines model parameters instead of samples. Other works have aimed to learn the SE(3) and periodic invariances through data augmentations [104, 31] or data canonicalization [101] with large language models (LLMs) [31, 101] or with diffusion of raw attributes [104] or of latents [42]. Recent works have used LLMs to generate noisy crystals or textual context which is passed as input to graph-based diffusion or flow matching [86, 103, 71]. We view many of these existing works as orthogonal to ours since our handling of space group symmetries can be combined with their modeling frameworks.

Two of the aforementioned works attempted to learn space group-conditioned generation without hard constraints. The graph diffusion model MatterGen [107] was fine-tuned on 14 space groups and used ground truth numbers of atoms per unit cell per space group to initialize generation. However, they could only generate stable, unique, and novel crystals in target space groups with 16% accuracy. Similarly, CrystalLLM [31] only managed 24% accuracy despite a generous SpaceGroupAnalyzer [91, 70] tolerance of 0.2Å.

Relevant to our model, a growing number of works have considered hard space group constraints during generation; however, to our knowledge, they either produce crystals without space group invariant likelihoods, thus assigning different likelihoods to symmetrically equivalent atoms, or rely on local structure relaxations using machine learning interatomic potentials (MLIPs) or DFT. WyCryst [108] trained a VAE to generate elements and Wyckoff position occupations, but rely on expensive DFT relaxations of atoms from uniform random locations in the Wyckoff positions. Crystal-GFN [64] considered space group constraints for the task of distribution matching under the GFlowNet framework [6] but did not address how to sample atom coordinates. DiffCSP++ [41] trained a graph-based diffusion model with masked diffusion of the unit cell lattice, continuous element diffusion with a post-hoc `argmax`, and projected diffusion of atom positions on the Wyckoff subspaces. However, DiffCSP++ is not space group equivariant to the best of our knowledge and relies on fixed templates from the training data of space group, number of atoms per unit cell, and occupied Wyckoff positions. SymmCD [55] is a non-equivariant diffusion model which leverages asymmetric units to reduce memory footprints; they use discrete diffusion of Wyckoff positions and elements and post-hoc projections of atomic coordinates onto Wyckoff positions. SymmBNF [78] recently extended SymmCD to the Bayesian Flow Networks framework [27], also requiring post-hoc projections of atom positions to the Wyckoff subspaces. CrystalFormer [10] trained an autoregressive transformer, canonicalizing crystals as a sequence of atoms ordered lexicographically by Wyckoff letter and then fractional coordinates. However, they rely on von Mises distributions which are not space group invariant. WyckoffTransformer [45] also autoregressively predicts atom types and Wyckoff positions but relies on DiffCSP++ or MLIPs to refine atoms from uniform random coordinates in the Wyckoff positions. We note that MLIPs are hindered by kinetic barriers and only preserve space group symmetry if they predict conservative forces of an invariant energy, generally requiring an expensive backward pass for inference [24, 77, 51]. WyckoffDiff [47] generates Wyckoff positions and elements with D3PM-based diffusion [3] but similarly relies on MLIPs to determine atom coordinates. Unlike these existing works, our model learns to generate complete crystals from scratch and produces space group invariant likelihoods via space group equivariant diffusion.

3 Preliminaries

Space groups Formally, a space group $G \in \mathbb{G}$, where \mathbb{G} denotes the set of 230 space groups, is a group of isometries that tiles \mathbb{R}^3 . In particular, G is generated by an infinite subgroup of discrete lattice translations $T_L = \{n_1 L_1, n_2 L_2, n_3 L_3 | n_i \in \mathbb{Z}, L_i \in \mathbb{R}^3\}$ as well as a collection of other symmetry operations $g(\cdot) = \{R_g(\cdot) + v_g | R_g \in O(3), v_g \in \mathbb{R}^3\} \in G$, where R_g is a point group operation (rotation, reflection, or identity) and v_g is a translation.

Wyckoff positions Given a space group and a point $x \in \mathbb{R}^3$, the *stabilizer group* $G_x \equiv \{g | gx = x\} \subset G$ (also called the site symmetry group in materials science) is the finite subgroup of G that leaves x invariant. A Wyckoff position is then defined as the set of points with conjugate stabilizer groups, i.e., $\{x' | \exists g \in G : G_{x'} = gG_xg^{-1}\}$. Conceptually, if g is a point group operation, this means that all points in a Wyckoff position are invariant to the same space group operations up to a change

of basis. By convention, when x is described with respect to the lattice basis $\{L_1, L_2, L_3\}$, the size of the *orbit* of x in the unit cell, $|\{gx|g \in G, gx \in [0, 1]^3\}|$, is called the *Wyckoff multiplicity*. Wyckoff positions whose stabilizer groups are non-trivial, i.e., include more than the identity operation, are referred to as *special Wyckoff positions* as opposed to the *general Wyckoff position* defined by the identity stabilizer group. Wyckoff positions are labeled by multiplicity and *Wyckoff letter*, where the lexicographic ordering gives the Wyckoff positions in (partial) order of increasing multiplicity.

Unit cells The infinite translational periodicity of a crystal can be represented with a parallelepiped Γ called the *unit cell*. The unit cell reduces infinite crystals by removing redundancy induced by T_L , the group of discrete lattice translations. In this way, crystals are represented by the tuple $M = (A, X, L)$, where $A = (a'_1, \dots, a'_N) \in \mathbb{A}^N$ are the atom types, \mathbb{A} is the set of all chemical elements, and N is the number of atoms in the unit cell; $X = \{(x'_1, \dots, x'_N) | x'_i \in \Gamma\}$ are the Cartesian atom coordinates; and $L = (L_1, L_2, L_3) \in \mathbb{R}^{3 \times 3}$ are the unit cell basis vectors. Given M , the infinite periodic structure can be reconstructed by applying the actions of T_L as $\{(a'_i, x'_i + n_1 L_1 + n_2 L_2 + n_3 L_3)_{i=1}^N | n_j \in \mathbb{Z}\}$. Alternatively, the atom coordinates can be given in the lattice basis instead of the Cartesian basis. Coordinates in the lattice basis are commonly referred to as *fractional coordinates*. In fractional coordinates, the infinite crystal is reconstructed as $\{(a'_i, x'_i + n_1 e_1 + n_2 e_2 + n_3 e_3)_{i=1}^N | n_j \in \mathbb{Z}\}$. For the rest of this paper, we assume atom coordinates are always in fractional coordinates. While the choice of unit cell is not unique, prior crystal generative models [100, 40, 62] either canonicalize it with a minimum-volume *primitive cell* determined by the Niggli algorithm [29] or a *conventional cell* [2] which contains all the symmetries of the space group.

Asymmetric units In our work, we represented crystals with a convex polytope $\Pi \in \mathbb{R}^3$, the asymmetric unit (ASU), which maximally reduces infinite crystals by removing all redundancies induced by the space group $G \supseteq T_L$. Under this formulation, we consider atoms in the ASU with fractional coordinates $X = \{(x_1, \dots, x_n) | x_i \in \Pi\} \in \mathbb{R}^{n \times 3}$, atom types $A = (a_1, \dots, a_n) \in \mathbb{A}^n$, and Wyckoff positions $W = \{(w_1, \dots, w_n) | w_i = G_{x_i}\} \in \mathbb{W}^n$ where $n \leq N$. The infinite periodic structure of a crystal can be reconstructed by applying the actions of G to Π , i.e.,

$$\{(a_i, g_{ij} x_i) | x_i \in \Pi, g_{ij} \in G/w_i, i \in (1, \dots, n)\}.$$

By focusing on these n symmetrically inequivalent atoms, we reduced our model's memory footprint and minimized the dimensionality of the generative modeling task. We canonicalized the non-unique choice of ASU using those listed in the International Tables for Crystallography [2] with additional conditions on faces, edges, and vertices from Grosse-Kunstleve et al. [30] to ensure that the ASUs are *exact*, i.e., that Π tiles \mathbb{R}^3 without overlaps at the boundaries $\partial\Pi$.

Diffusion models Diffusion-based generative models [33, 84] construct samples $\mathbf{x}_0 \sim p_{t=0}$ from a noisy prior $\mathbf{x}_T \sim p_{t=T}$ by evolving \mathbf{x}_T through time t with a learned denoising process. The training process generally involves iteratively adding noise to data samples \mathbf{x}_0 via the forward stochastic differential equation (SDE),

$$d\mathbf{x} = \mathbf{f}(\mathbf{x}, t)dt + \eta(t)d\mathbf{w}, \quad (1)$$

where $\mathbf{f}(\mathbf{x}, t) : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is the drift coefficient, \mathbf{w} is the standard Wiener process, and $\eta(t) \in \mathbb{R}$ is the diffusion coefficient. Samples are generated by running the process backwards in time from T to 0 via the *reverse* SDE,

$$d\mathbf{x} = [\mathbf{f}(\mathbf{x}, t) - \eta(t)^2 \nabla_{\mathbf{x}} \log p_t(\mathbf{x})]dt + \eta(t)d\bar{\mathbf{w}}, \quad (2)$$

where $\bar{\mathbf{w}}$ is another standard Wiener process. The generative model aims to learn the score of the marginal distribution $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$, and then simulate Eq. 2 to sample from p_0 . In our model, we have $\mathbf{f}(\mathbf{x}, t) = 0$.

Space group equivariant functions For a group of isometries G' , Rezende et al. [76] and Köhler et al. [52] showed that equivariant flows preserve invariant probability densities. Conversely, the gradient of a G' -invariant function is equivariant [44]. We are interested in using this result for the space groups. A function $f : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ is space group equivariant if it satisfies $R_g f(x) = f(gx)$ for all $g \equiv (R_g, v_g) \in G$, where $R_g \in O(3)$ is a point group element and $v_g \in \mathbb{R}^3$ is a translation. *Symmetrization* via summing over group elements has become a popular method to build invariant and equivariant neural networks in the last few years [105, 74]. Recently, Miramezani et al. [63]

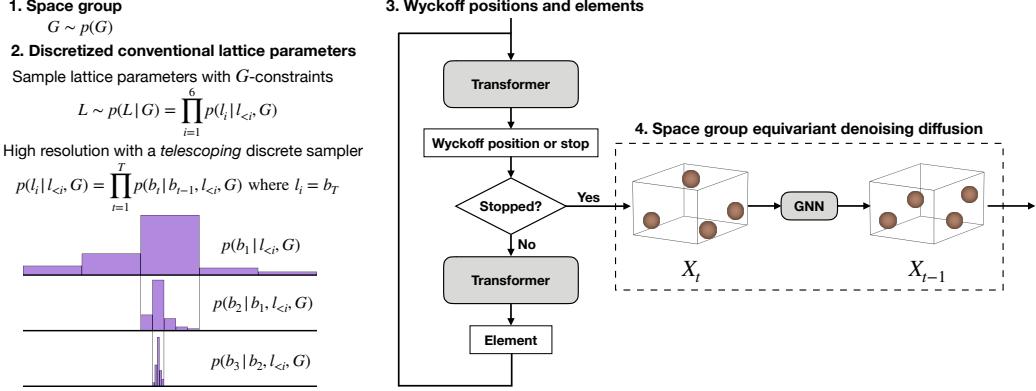


Figure 2: Illustration of our crystal generation process.

showed how symmetrization can be applied to space groups despite their countably infinite size. In particular, given the finite quotient group G/T_L and a T_L -invariant function $\hat{f} : \mathbb{R}^3 \rightarrow \mathbb{R}^3$, the following function is space group equivariant:

$$f(x) = \frac{1}{|G/T_L|} \sum_{g \in G/T_L} R_g^{-1} \hat{f}(gx). \quad (3)$$

Besides faithfully modeling the space group symmetries of crystals, we discovered that space group equivariant vector fields conveniently live in the tangent spaces of the Wyckoff positions. We formalize this in the following theorem.

Theorem 1 Let G be a space group, $x \in \mathbb{R}^3$ be a point residing in a Wyckoff position with stabilizer group G_x , and $f : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ be a space group equivariant function. Then, for arbitrary constant $c \in \mathbb{R}$ and any space group element $g \in G_x$,

$$g(x + cf(x)) = x + cf(x) \implies G_x \subseteq G_{x+cf(x)}. \quad (4)$$

See proof in Sec. A.1. This result implies that a space group equivariant flow is automatically restricted to the manifold of its Wyckoff position, preventing the need for projections that lead to indeterminate probability distributions [41] or a discontinuous generation process [55, 78]. The result also implies that if, for example, a 0D Wyckoff position (a point) lies in a 1D Wyckoff position (a line), then a space group equivariant vector field cannot move an atom along the line through the point. This is problematic because we observed examples where traversing between symmetrically unique regions of a 1D Wyckoff position requires moving through a 0D Wyckoff position (e.g., in space group 192, the 1D Wyckoff position 12k is split by the 0D Wyckoff position 4c. See Fig. 3). This issue motivates SDE-based sampling since Gaussian noise (projected onto the tangent space of the Wyckoff position) can stochastically perturb atoms through 0D Wyckoff positions while maintaining that the marginal distribution over atom coordinates is still space group invariant.

4 Proposed method: SGEquiDiff

4.1 Crystal representation

Using `sympy` [61] and `PyXtal` [22], we removed redundancy induced by space group symmetry from every Wyckoff position by intersecting each one with its exact asymmetric unit. Each zero-dimensional Wyckoff position was reduced to a single point, each one-dimensional position to a set of line segments, each two-dimensional position to a set of convex polygonal ASU facets, and each three-dimensional position to the ASU interior. We will refer to these convex sets as *Wyckoff shapes*. Intersections for space group 192 are shown as an example in Figure 3. We used Wyckoff shapes for memory-efficient crystal representation and sampling as will be described in the following sections.

We decomposed each crystal into $M = (G, L, A, W, X) \in (\mathbb{G}, \mathbb{R}^{3 \times 3}, \mathbb{A}^n, \mathbb{W}^n, \mathbb{R}^{n \times 3})$, where n is the number of atoms in the asymmetric unit, $G \in \mathbb{G}$ is the space group, $L \in \mathbb{R}^{3 \times 3}$ is the

conventional lattice basis, $A \in \mathbb{A}^n$ are elements, $W \in \mathbb{W}^n$ are Wyckoff positions, and $X \in \mathbb{R}^{n \times 3}$ are fractional atom coordinates. Several variables in M impose hard constraints on each other. To handle these dependencies, the model sequentially samples each constrained variable after its constraining variable(s). We factorized the generation process as follows:

$$p(M) = p(G) \times p(L|G) \times p(W, A|L, G) \times p(X|W, A, L, G). \quad (5)$$

Structuring the generation process in this way rather than jointly modeling all crystal components avoids tuning per-component loss weights [100], memory-intensive mask states [104], and post-hoc projections that are intractable to marginalize over [55]. We set $p(G)$ to the empirical training distribution. Now we describe our method for the remaining conditional distributions.

4.2 Telescoping discrete sampling of L

We parameterized univariate conditionals to autoregressively sample the 3 lattice lengths (a, b, c) and 3 angles between them (α, β, γ) for conventional unit cells, applying physical constraints to each conditional. Canonicalizing the lattice with the conventional lattice parameters makes the lattice sampling automatically SE(3) invariant. Denoting $l = (a, b, c, \alpha, \beta, \gamma)$, the model learned

$$p(L|G) = \prod_{i=1}^6 p(l_i|l_{<i}, G), \quad (6)$$

where $p(l_i|l_{<i}, G)$ has support over positive values with finite range determined by the data. Under the space group constraints, crystal lattices can be binned into 6 crystal families, each putting unique constraints on the lattice parameters. Space groups 1 to 2 impose no constraints; 3 to 15 require $\alpha = \gamma = 90^\circ$; 16 to 74 require $\alpha = \beta = \gamma = 90^\circ$; 75 to 142 require $\alpha = \beta = \gamma = 90^\circ$ and $a = b$; 143 to 194 require $a = b$, $\alpha = \beta = 90^\circ$, and $\gamma = 120^\circ$; and 195 to 230 require $a = b = c$ and $\alpha = \beta = \gamma = 90^\circ$. Furthermore, the crystal lattice must have non-zero volume. To impose these constraints, the model only learns univariate conditionals for the lattice parameters unconstrained by the crystal families, leaving constrained terms in the product of Equation 6 equal to 1. We enforce positive volume by dynamically setting the support of $p(\gamma|a, b, c, \alpha, \beta, G)$ to satisfy

$$\frac{\text{Volume}}{abc} = \sqrt{1 + 2 \cos \alpha \cos \beta \cos \gamma - \cos^2 \alpha - \cos^2 \beta - \cos^2 \gamma} > 0.$$

Besides physical constraints, lattice generation requires the flexibility to learn highly peaked distributions since small perturbations to a crystal lattice can significantly alter materials properties. In BaTiO₃ for example, 0.03 Å strain was found to increase the ferroelectric transition temperature by 500°C and the remnant polarization by 250% [13]. In contrast, the range of conventional lattice lengths in the MP20 dataset is over 100 Å. To address this challenge, we chose to discretize the lattice parameters to a resolution of 0.01 Å. Naively, this resolution requires a softmax over $N_l = \mathcal{O}(10^4)$ classes per lattice parameter to achieve a 100 Å range. We overcame this poor scaling by *telescoping* the categorical distribution. See Figure 2 for a visual explanation. At a high level, the range of lattice parameters was first binned very coarsely, and a class b_1 was sampled from $p(b_1|l_{<i}, G)$. Then, the selected class b_1 was further coarsely binned and one of these higher resolution bins was selected from $p(b_2|b_1, l_{<i}, G)$. This process was repeated K times to achieve higher levels of resolution as

$$p(l_i|l_{<i}, G) = \prod_{k=1}^K p(b_k|b_{k-1}, l_{<i}, G)$$

where $l_i = b_K$ and $b_0 = \emptyset$. Choosing $p(b_k|b_{k-1}, l_{<i}, G)$ to be a categorical distribution over a small number of classes $n_l \ll N_l$ achieves $(1/n_l^K) = (1/N_l)$ resolution with $\mathcal{O}(n_l K) \ll \mathcal{O}(N_l)$ memory. We used $K = 2$ and $n_l = 100$ in our experiments. We represented conditioning on l_i and b_k by using min-max normalization and Gaussian random Fourier features [89]. Logits for $p(b_k|b_{k-1}, l_{<i}, G)$ were produced with an MLP. Space group constraints were enforced with hard-coded masking of lattice parameters and their logits.

To better align the distributions of partially complete lattices seen during training and inference, we employed noisy teacher forcing during training. Specifically, we minimized the negative log likelihood of the next ground truth lattice parameter conditioned on noisy versions of previous ground truth lattice parameters, where noise was masked to respect space group constraints. Lattice lengths and angles were augmented with uniform random noise with 0.3 Å and 5° ranges, respectively. Noising was not applied at inference time.

4.3 Transformer-based sampling of W , A , and n

Similarly to Cao et al. [10], Kazeev et al. [46], and Kelvinius et al. [47], we leveraged a transformer architecture [92] to autoregressively predict atomic Wyckoff positions $W \in \mathbb{W}^n$ and elemental species $A \in \mathbb{A}^n$ as

$$p(W, A|L, G) = \prod_{i=1}^n \left[p(w_i|w_{<i}, a_{<i}, L, G) p(a_i|w_{\leq i}, a_{<i}, L, G) p(\text{stop}|w_{\leq i}, a_{\leq i}, L, G) \right]. \quad (7)$$

The number of atoms in the ASU n is sampled implicitly via $p(\text{stop}|w_{\leq i}, a_{\leq i}, L, G)$. We used attention masking to enforce unique token orderings per crystal during training and inference. Specifically, atoms were ordered lexicographically first by Wyckoff letter and then by atomic number. The transformer input is invariant to permutations of any remaining ties. We additionally set logits of zero-dimensional Wyckoff positions which are already occupied to negative infinity to prevent sampling overlapping atoms at those locations. The transformer was fit with the cross-entropy loss. See section A.4 for details of our transformer architecture. We did not use a positional encoding since the model does not need to distinguish between different orderings of the same atoms.

To aid generalization across the conditioning variables, we initialized embeddings from physical descriptors. Element embeddings were initialized as those introduced by Xie and Grossman [99]. Similarly to the lattice sampler, min-max normalized lattice parameters were embedded with Gaussian random Fourier features. Learnable stop tokens were initialized from Gaussian noise. We represented each of the 230 space groups with 62 initial features and each of the 1731 Wyckoff positions across space groups with 231 initial features. For details and an ablation against the symmetry features used in SymmCD [55], see Sec. A.4.3.

4.4 Space group equivariant diffusion of X

We leveraged the score matching framework [84, 36] to generate atomic fractional coordinates X . Unlike prior works that fit to scores of the translation-invariant Wrapped Normal (WN) distribution [107, 41, 40], we learned scores of a G -invariant *Space Group Wrapped Normal* (SGWN) distribution:

$$p(x_t|x_0) \propto \sum_{g \in G} \exp\left(\frac{-\|x_t - gx_0\|_F^2}{2\sigma_t^2}\right). \quad (8)$$

Theorem 2 The SGWN is G -invariant, i.e., $\forall g = (R_g, v_g) \in G$, $p(gx_t|x_0) = p(x_t|x_0)$. Additionally, the score is space group equivariant, i.e., $\nabla_{x_t} \log p(R_g x_t + v_g|x_0) = R_g \nabla_{x_t} \log p(x_t|x_0)$ (see Sec. A.1 for proof).

Following Jiao et al. [40], we set the noise scale σ_t with the exponential scheduler: $\sigma_0 = 0$ and $\sigma_t = \sigma_1 (\frac{\sigma_T}{\sigma_1})^{\frac{t-1}{T-1}}$ if $t > 0$. As σ_t gets large, the SGWN approaches the uniform distribution. Noisy samples x_t were created by sampling Gaussian noise $\epsilon \sim \mathcal{N}(0, I)$ and reparameterizing as $x_t = x_0 + \mathcal{P}_{x_0}(\sigma_t \epsilon)$ where $\mathcal{P}_{x_0}(\cdot)$ is an orthogonal projection to the tangent space of x_0 's Wyckoff shape. For the backward process, we constructed the uniform prior $p(x_T)$ on each Wyckoff position. To do so without redundantly assigning probability mass to symmetrically equivalent points, we placed uniform distributions on the relevant Wyckoff shapes. See Sec. A.3 for details. We used predictor-corrector sampling [85] to sample x_0 .

To learn the space group equivariant scores of the SGWN distribution, we built a space group equivariant graph neural network s_θ (see Sec. A.4 for architectural details) and trained it with the following score matching loss:

$$L_X = \mathbb{E}_{x_t \sim p_t(x_t|x_0)p(x_0), t \sim \mathcal{U}(0, T)} [\|s_\theta(x_t, t) - \lambda_t \nabla_{x_t} \log p(x_t|x_0)\|_2^2]$$

where $\lambda_t = \mathbb{E}[\|\nabla_{x_t} \log p(x_t|x_0)\|]$ was employed for training stability. Approximate computation of λ_t and $\nabla_{x_t} \log p(x_t|x_0)$ is described in Sec. A.2.

To save memory when symmetrizing s_θ with Equation 3, we note that only $|G|/|G_x| \leq |G|/|T_L|$ elements in the set of points $\{gx|g \in G/T_L\}$ are unique. Thus we only predicted $\hat{f}_\theta(gx)$ on unique points and subsequently indexed into these predictions to compute the full sum over the $|G|/|T_L|$ summands in Equation 3. For points $x \in \mathbb{R}^3$ in special Wyckoff positions with large stabilizer groups G_x , this significantly reduced the number of forward passes through the model.

Table 1: **MP-20 dataset.** We report metrics of validity for 10,000 sampled crystals; uniqueness and novelty for 1,000 valid crystals; diversity and goodness-of-fit distributional distances for valid, unique, and novel crystals; and stability for valid, unique, and novel crystals. Stability is defined as having a DFT-predicted $E_{\text{hull}} < 0.1$ eV/atom with respect to the Materials Project v2022.8.23. T represents the inference-time temperature of discrete distributions in SGEquiDiff. Unless noted otherwise, $T=1.0$. Best result is bolded, second best is underlined. Standard deviations over three training runs are reported in parentheses as errors of rightmost digits.

	Time ↓ (s / batch)	Validity (%) ↑ Structure	Composition	U.N. rate (%) ↑ Template	Structure	W_p	Distribution distance ↓ $W_{N_{el}}$	JSD_G	$JSD_{d_{\text{Wyckoff}}}$	CMD ↓ Structure	Diversity ↑ Structure	Composition	S.U.N. (%) ↑ MP-2024
CDVAE [100]	906	99.99	85.66	15.5	98.2	0.6590	1.423	0.6957	0.4590	0.4821	0.6539	13.70	14.2
DiffCSP [40]	<u>154</u>	99.92	82.21	12.2	85.6	0.1454	0.4000	0.4638	0.2328	0.1766	<u>0.9588</u>	15.69	-
DiffCSP++ [41]	484	99.92	85.94	1.1	84.7	<u>0.1658</u>	0.5002	<u>0.1608</u> [*]	0.0449 [*]	0.1079	0.9329	15.23	23.4
SymmCD [55]	139	88.24	86.76	9.6	87.7	<u>0.1640</u>	0.3213	<u>0.1669</u>	0.0344	0.3233	0.9111	15.62	-
FlowMM [†] [62]	-	96.85	83.19	-	-	-	-	-	-	-	-	-	-
MatterGen [107]	3,689	100.0	82.6	15.2	86.3	0.2059	0.2416	0.4331	0.2129	<u>0.1338</u>	0.9889	15.67	<u>24.3</u>
SGEquiDiff	226 ₍₁₂₎	99.25 ₍₂₈₎	<u>86.16</u> ₍₃₂₎	18.6 ₍₁₆₎	80.0 ₍₁₆₎	0.1939 ₍₁₀₄₉₎	0.2092 ₍₆₇₈₎	<u>0.1733</u> ₍₅₃₎	0.0252 ₍₉₀₎	0.1718 ₍₅₈₎	0.9190 ₍₂₀₄₎	15.64 ₍₇₎	24.4
SGEquiDiff ($T=1.5$)	-	98.97 ₍₁₈₎	83.36 ₍₅₂₎	31.9 ₍₁₅₎	90.1 ₍₁₄₎	0.3810 ₍₁₄₄₀₎	0.3009 ₍₁₄₈₎	0.2183 ₍₇₃₎	0.0436 ₍₄₃₎	0.2932 ₍₉₂₎	0.8556 ₍₁₀₁₎	15.82 ₍₁₈₎	-
SGEquiDiff ($T=2.0$)	-	98.40 ₍₃₇₎	81.08 ₍₂₂₎	42.8 ₍₃₁₎	<u>95.0</u> ₍₂₎	0.6524 ₍₉₃₇₎	0.3685 ₍₄₀₉₎	0.2698 ₍₁₀₈₎	0.0564 ₍₆₁₎	0.3760 ₍₁₁₅₎	0.8439 ₍₁₃₇₎	16.28 ₍₁₇₎	-
SGEquiDiff ($T=3.0$)	-	97.84 ₍₅₃₎	78.60 ₍₃₅₎	57.1 ₍₃₂₎	<u>95.0</u> ₍₅₎	1.473 ₍₂₆₇₎	0.4962 ₍₁₇₇₎	0.3367 ₍₂₀₄₎	0.0751 ₍₄₁₎	0.4656 ₍₁₃₂₎	0.8177 ₍₁₀₁₎	17.16 ₍₃₀₎	-

^{*} Uses fixed templates from the training data. [†] Values reported by Miller et al. [62].

5 Experiments

Evaluations. We evaluated SGEquiDiff on two benchmark datasets: MP20 [100], containing 45,231 experimentally known crystals with up to 20 atoms per unit cell, and the more challenging MPTS52 dataset [4], containing 40,476 experimentally known crystals with up to 52 atoms per unit cell. Evaluations were conducted on 10,000 generated crystals. Following Xie et al. [100], we computed structural and compositional validity percentages using heuristics of interatomic distances and charge, respectively. We note that only $\sim 90\%$ of crystals in MP20 pass the composition validity checker based on SMACT [16]. Of the 10,000 generated crystals, we randomly sampled 1,000 determined to be both structurally and compositionally valid. Of these 1,000 crystals, we determined how many were unique and novel with respect to the training data using pymatgen’s StructureMatcher [70] with $\text{stol}=0.3$, $\text{angle_tol}=5$, and $\text{ltol}=0.2$ (*U.N. structures*). We also followed SymmCD and calculated the fraction of unique and novel *templates* (*U.N. templates*), where a template is defined as a space group and multiset of occupied Wyckoff positions. U.N. structures were used to compute (1) distribution distances between ground truth test and generated materials properties, including Wasserstein distances for atomic density ρ and number of unique elements N_{el} as well as Jensen-Shannon divergences (JSD) for space group G and occupied Wyckoff dimensionalities d_{Wyckoff} ; (2) Central Moment Discrepancy (CMD) [106] up to 50 moments between ground truth test and generated CrystalNN structural fingerprints [109]; and (3) structural and compositional diversity as measured by average pairwise L_2 -distances between CrystalNN and Magpie [93] fingerprints, respectively. Average sampling times per batch of 500 crystals were measured on an NVIDIA A40 GPU. We list hyperparameters and training times for SGEquiDiff in Sec. A.7.

Finally, for all U.N. structures out of 1,000 random valid crystals, we conducted structure relaxations with expensive density functional theory (DFT) calculations to assess thermodynamic stability. Aligned with analyses of DFT-calculated energies for experimentally observed crystals [88] and of DFT errors relative to experiment [87, 5], stability was defined as having a predicted *energy above the hull* less than 0.1 eV/atom with respect to crystals in the Materials Project v2022.8.23 [39]. Further details on the DFT calculations are in Sec. A.8. The fraction of stable materials was multiplied by the fraction of U.N. structures to calculate the stable, unique, and novel (S.U.N.) rate. We qualify that a low energy above the hull as predicted by DFT has been found to be a necessary but insufficient condition for synthesizability [88]; developing physical theories for predictive synthesis is an active area of research in the materials science community [60, 25, 43, 12, 68, 48]. Due to the computational expense of DFT calculations, we only computed S.U.N. rates on a subset of baseline models. While prior works pre-relax generated crystals with foundational machine learning interatomic potentials (MLIPs) before conducting DFT relaxations [62, 86, 107], we opted not to use MLIPs to avoid conflating their biases [17, 53] with our evaluations of the generative models.

Baselines. We compared our model to several prior methods. CDVAE (4.9M parameters) [100] predicts lattices and numbers of atoms per unit cell with a VAE and then samples elements and coordinates with denoising diffusion. DiffCSP (12.3M parameters) [40] jointly diffuses the lattice with the atoms using fractional atom coordinates. MatterGen (44.6M parameters) [107] similarly diffuses the lattice, atom types, and atom positions, but leverages Cartesian atom coordinates. FlowMM [62] extends DiffCSP with the flow matching framework. DiffCSP++ (12.3M parameters) [41] extends DiffCSP by enforcing space group constraints with projected diffusion of the lattice and

Table 2: **MPTS-52 dataset.** We report metrics of validity for 10,000 sampled crystals; uniqueness and novelty for 1,000 valid crystals; and diversity and distributional distances for valid, unique, and novel crystals. Best result is bolded, second best is underlined. For DiffCSP and DiffCSP++, we used the same diffusion corrector step sizes as for MP-20, and the DiffCSP++ batch size was decreased to 96 to avoid OOMs. CDVAE failed to train due to graph construction errors from isolated atoms. Standard deviations over three training runs are reported in parentheses as errors of rightmost digits.

	Time ↓ (s / batch)	Validity (%) ↑ Structure	U.N. rate (%) ↑ Composition	Template Structure	W_p	Distribution distance ↓ $W_{N_{el}}$	JSD_G	$JSD_{d_{Wyckoff}}$	CMD ↓ Structure	Diversity ↑ Structure	Composition
DiffCSP [40]	467	67.47	55.8	19.8	79.8	1.189	0.5006	0.6900	0.0818	0.4836	0.8621
DiffCSP++ [41]	1230	99.87	77.52	1.0	86.9	0.8244	0.3692	0.2759	0.1154	0.3812	0.8457
SymmCD [55]	210	87.11	<u>78.18</u>	14.8	90.2	1.126	0.3506	<u>0.2730</u>	0.1146	0.4775	0.7843
SGEquiDiff	530 ₍₄₀₎	<u>97.79₍₃₇₎</u>	79.83₍₉₆₎	<u>38.7₍₂₅₎</u>	89.8₍₇₎	0.6110₍₁₅₉₂₎	<u>0.1736₍₄₂₈₎</u>	0.3104 ₍₆₂₎	0.1300 ₍₁₄₂₎	<u>0.4186₍₂₈₎</u>	0.8613₍₂₉₂₎
											<u>16.31₍₁₁₎</u>

atom coordinates; however, they do so without space group equivariant scores. SymmCD (60.4M parameters) [55] enforces space group constraints with non-equivariant diffusion of atoms in the asymmetric unit followed by post-hoc projections to the Wyckoff positions.

Results. On MP20, we found SGEquiDiff was competitive on all metrics and achieved the highest S.U.N. rate (Table 1). While MatterGen achieved a comparable S.U.N. rate, we note that it is a significantly larger model (44.6M parameters) than SGEquiDiff (5.5M parameters), sampled crystals \sim 16 times slower, and yielded worse JSD metrics for space groups and Wyckoff dimensions. Compared to DiffCSP++, SGEquiDiff achieved a slightly higher S.U.N. rate and significantly higher U.N. template rate with \sim 2.2 times fewer model parameters. Unsurprisingly, SGEquiDiff and other space group-constrained models (DiffCSP++, SymmCD) strongly outperformed baselines on the JSD metric for space groups and Wyckoff dimensions. We also show that by increasing the inference-time temperature of SGEquiDiff’s categorical distributions over space group, Wyckoff position, element, and the stop token, the U.N. rates and composition diversity metric can be significantly improved at the expense of slightly lower validity rates and higher distributional distances to the test set. Inspecting normalized root mean square Cartesian displacements between as-generated and DFT-relaxed structures, we found all models performed similarly, with MatterGen being the best on average (Table 4). Examples of S.U.N. crystals from SGEquiDiff are shown in Fig. 4.

We also benchmarked SGEquiDiff on the more difficult MPTS52 dataset (Table 2), which most other models do not present results on. Using the same hyperparameters as used on MP20, we found that SGEquiDiff was able to scale to more atoms while maintaining a high rate of valid, unique, and novel crystals. SGEquiDiff notably achieved the highest U.N. template rate, composition validity, and Wasserstein distances for density and number of elements. Interestingly, all space group-constrained models significantly outperformed DiffCSP in the validity metrics, highlighting the advantage of enforcing nontrivial space group symmetries as an inductive bias.

Ablations. To isolate effects of various components of SGEquiDiff, we performed ablations on the MP-20 dataset. First, we replaced our autoregressive lattice sampler with the DDPM-based [33] sampling of E(3)-invariant lattice matrix representations proposed in DiffCSP++ and used in SymmCD (Table 6). We used the same cosine noise scheduler as DiffCSP++, jointly diffusing lattices and atom coordinates. We term this variant of SGEquiDiff as *+LDiff*. We observed that *+LDiff* improved the Wasserstein distance for number of elements (possibly because conditioning of the Wyckoff-Element transformer on lattices was removed) and CMD of structural fingerprints. However, our autoregressive lattice sampler yielded better validity, JSD of Wyckoff dimensionalities, and S.U.N. rate. We also observed that in roughly 1% of samples, the *+LDiff* model produced lattice matrices with negative determinant, resulting in a spurious inversion that does not preserve the symmetry of the 22 chiral space groups. We show examples of S.U.N. crystals from *+LDiff* in Fig. 5

We also replaced our space group and Wyckoff position features with those introduced in SymmCD [55] (Table 5). We found our featurizations yielded better structure validity, composition validity, Wasserstein distance for atomic density, structural divergence, and validation log-likelihoods of ground truth Wyckoff positions. In contrast, SymmCD features yielded more diverse samples, evidenced by higher U.N. rates and composition diversity.

Since space group symmetrization via Eq. 3 can be applied to any periodic translation invariant model, we replaced our FAENet-inspired [20] graph neural network (GNN) with the EGNN-inspired [80] model (*CSPNet*) used by other crystal generative models [40, 41, 62]. Under a fixed number of trainable parameters (2.2M), we found our GNN yielded better metrics for structural validity, composition validity, and CMD of structural fingerprints (Table 5).

Crystal structure prediction (CSP). Consistent with previous works [40, 41, 62], we evaluated SGEquiDiff on its ability to predict crystals from a test set using elemental composition as input. Performance was assessed by calculating match rate (MR) and root mean squared error (RMSE) of atomic fractional coordinates based on a single generated crystal per composition. RMSE was normalized by the average free length per atom. Matching was done with `pymatgen StructureMatcher` [70] using `stol=0.5`, `angle_tol=10`, and `1tol=0.3`. To adapt SGEquiDiff to this task, we combined the pretrained `+LDiff` variant of SGEquiDiff with the metric learning-based approach proposed in DiffCSP++ (*CSPML*) for selecting templates from training data and performing element substitutions to satisfy the given composition. Without CSP-specific training (unlike baselines), SGEquiDiff+LDiff achieved competitive performance (Table 3). We note SGEquiDiff could alternatively be adapted for CSP by reordering the factorization in Eq. 5 or by supplying composition as conditioning during training. While more challenging, learning instead of strictly enforcing composition may be desirable since an arbitrary composition is not guaranteed to host a stable crystal (accordingly, experimentally observed compositions can be weighted averages of multiple phases).

We highlight that the current CSP task is slightly misaligned with real-world application. Many existing models assume the number of atoms per unit cell will be known beforehand, which is not necessarily true in practice. Additionally, a single composition can host multiple stable crystal structures depending on experimental conditions. For example, iron can exist as a body-centered cubic (2 atoms per conventional cell), face-centered cubic (4 atoms per conventional cell), or at high pressures, hexagonal close-packed structure (6 atoms per conventional cell). Furthermore, materials discovery campaigns are often aimed at composition spaces with sparse training data, violating the i.i.d. assumption of the MP20 dataset splits [100].

Table 3: **Crystal structure prediction.** Leveraging the template selection and element substitution approach developed in DiffCSP++ [41], we evaluated SGEquiDiff+LDiff on the crystal structure prediction task introduced by DiffCSP [40].

MP-20	MR (%) \uparrow	RMSE \downarrow
CDVAE* [100]	33.90	0.1045
DiffCSP* [40]	51.49	0.0631
FlowMM* [62]	61.39	0.0566
DiffCSP++ (+CSPML) [41]	70.58	0.0272
SGEquiDiff (+LDiff,CSPML)	69.42	0.0416

* Uses ground truth number of atoms per unit cell.

6 Conclusion

In this paper, we proposed SGEquiDiff to generate crystals with space group invariant likelihoods by leveraging equivariant diffusion. We showed that space group equivariant flows automatically live on the manifolds of the Wyckoff positions. Significantly, SGEquiDiff achieved state-of-the-art generation rates of stable, unique, and novel crystals as evaluated by rigorous quantum mechanical simulations. We proposed an efficient autoregressive method for sampling space group-constrained lattices. We also showed that tuning the inference-time temperature for autoregressive sampling of discrete crystal attributes provides control over the novelty of generated crystals.

Limitations and future directions. Due to computational constraints, we were only able to calculate S.U.N. rates on 1,000 valid crystals per model, adding unknown variance to reported metrics. SGEquiDiff’s training data was pre-processed to assign space groups to crystals with nonzero tolerances over atom positions and lattice angles; it is possible that some materials were incorrectly symmetrized to higher symmetry space groups than representative of reality. SGEquiDiff also does not provide guidance for how to synthesize generated materials. Other symmetry groups relevant to crystals were ignored, including magnetic space groups [94], spin space groups [98], and layer groups [23]. Practically relevant deviations of real materials from perfect crystals were not considered, including defects, compositional disorder, surface effects, and interfacial effects. Future work might include application to a broader set of tasks such as those derived from non-scalar properties [102, 59, 95] or materials spectroscopy [56]. Furthermore, training non-equivariant models with space group invariant target distributions and/or applying inference-time space group symmetrization may enable a broader diversity of model architectures to generate high symmetry crystals.

Broader impacts. This work has the potential to accelerate discovery of advanced materials for energy, electronics, optics, catalysis, aerospace, and more. Possible negative impacts include development of materials that are toxic, require energy-intensive processing, lead to depletion of raw minerals, or are used for military applications.

Acknowledgments and Disclosure of Funding

We acknowledge helpful discussions with Cindy Zhang, Barry Bradlyn, Alex M. Ganose, and Eric Toberer. This research used the Delta advanced computing and data resource (award OAC 2005572) and the Illinois Campus Cluster, operated by the Illinois Campus Cluster Program in conjunction with the National Center for Supercomputing Applications. The research was supported by the National Science Foundation under Grant Nos. DGE 21-46756 and 2118201.

References

- [1] Michael S. Albergo, Nicholas M. Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions, 2023. URL <https://arxiv.org/abs/2303.08797>.
- [2] M.I. Aroyo, H. Burzlaff, G. Chapuis, W. Fischer, H.D. Flack, A.M. Glazer, H. Grimmer, B. Gruber, Th. Hahn, H. Klapper, E. Koch, P. Konstantinov, V. Kopsky, D.B. Litvin, A. Looijenga-Vos, K. Momma, U. Mueller, U. Shmueli, B. Souvignier, J.C.H. Spence, P.M. de Wolff, H. Wondratschek, and H. Zimmerman. *International Tables for Crystallography: Space-group symmetry*, volume A. International Union of Crystallography, Chester, England, 2 edition, 2016. ISBN 978-0-470-97423-0. doi: 10.1107/97809553602060000114. URL <https://it.iucr.org/Ac/>.
- [3] Jacob Austin, Daniel D. Johnson, Jonathan Ho, Daniel Tarlow, and Rianne van den Berg. Structured denoising diffusion models in discrete state-spaces. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 17981–17993. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/958c530554f78bcd8e97125b70e6973d-Paper.pdf.
- [4] Sterling G. Baird, Hasan M. Sayeed, Joseph Montoya, and Taylor D. Sparks. matbench-
genmetrics: A python library for benchmarking crystal structure generative models using
time-based splits of materials project structures. *Journal of Open Source Software*, 9(97):5618,
2024. doi: 10.21105/joss.05618. URL <https://doi.org/10.21105/joss.05618>.
- [5] Christopher J. Bartel, Alan W. Weimer, Stephan Lany, Charles B. Musgrave, and Aaron M. Holder. The role of decomposition reactions in assessing first-principles predictions of solid stability. *npj Computational Materials*, 5(1):1–9, dec 2019. ISSN 20573960. doi: 10.1038/S41524-018-0143-2;SUBJMETA=1034,301,563,638,639;KWRD=THEORETICAL+CHEMISTRY,THEORY+AND+COMPUTATION. URL <https://www.nature.com/articles/s41524-018-0143-2>.
- [6] Emmanuel Bengio, Moksh Jain, Maksym Korablyov, Doina Precup, and Yoshua Bengio. Flow network based generative models for non-iterative diverse candidate generation. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021. URL <https://openreview.net/forum?id=Arn2E4IRjEB>.
- [7] G. Bergerhoff, R. Hundt, R. Sievers, and I. D. Brown. The Inorganic Crystal Structure Data Base. *Journal of Chemical Information and Computer Sciences*, 23(2):66–69, 1983. ISSN 00952338. doi: 10.1021/ci00038a003.
- [8] Peter E Blöchl. Projector augmented-wave method. *Physical review B*, 50(24):17953, 1994.
- [9] Jennifer Cano and Barry Bradlyn. Band Representations and Topological Quantum Chemistry. *Annual Review of Condensed Matter Physics*, 12:225–246, 2020. doi: 10.1146/annurev-conmatphys-041720-124134. URL <https://arxiv.org/abs/2006.04890><http://dx.doi.org/10.1146/annurev-conmatphys-041720-124134>.
- [10] Zhendong Cao, Xiaoshan Luo, Jian Lv, and Lei Wang. Space group informed transformer for crystalline materials generation, 2024.

[11] Lucas Caretta, Yu-Tsun Shao, Jia Yu, Antonio B Mei, Bastien F Grosso, Cheng Dai, Piush Behera, Daehun Lee, Margaret McCarter, Eric Parsonnet, Hari Krishnan K. P, Fei Xue, Xiang-wei Guo, Edward S Barnard, Steffen Ganschow, Zijian Hong, Archana Raja, Lane W Martin, Long-Qing Chen, Manfred Fiebig, Keji Lai, Nicola A Spaldin, David A Muller, Darrell G Schlom, and Ramamoorthy Ramesh. Non-volatile electric-field control of inversion symmetry. *Nature Materials*, 22(2):207–215, 2023. ISSN 1476-4660. doi: 10.1038/s41563-022-01412-0. URL <https://doi.org/10.1038/s41563-022-01412-0>.

[12] Jiadong Chen, Samuel R. Cross, Lincoln J. Miara, Jeong-Ju Cho, Yan Wang, and Wenhao Sun. Navigating phase diagram complexity to guide robotic inorganic materials synthesis. *Nature Synthesis*, 3(5):606–614, 2024. doi: 10.1038/s44160-024-00502-y. URL <https://doi.org/10.1038/s44160-024-00502-y>.

[13] K J Choi, M Biegalski, Y L Li, A Sharan, J Schubert, R Uecker, P Reiche, Y B Chen, X Q Pan, V Gopalan, L.-Q. Chen, D G Schlom, and C B Eom. Enhancement of Ferroelectricity in Strained BaTiO₃ Thin Films. *Science*, 306(5698):1005–1009, 2004. doi: 10.1126/science.1103218. URL <https://www.science.org/doi/abs/10.1126/science.1103218>.

[14] T Choi, S Lee, Y J Choi, V Kiryukhin, and S.-W. Cheong. Switchable Ferroelectric Diode and Photovoltaic Effect in BiFeO₃. *Science*, 324(5923):63–66, 2009. doi: 10.1126/science.1168636. URL <https://www.science.org/doi/abs/10.1126/science.1168636>.

[15] Stefano Curtarolo, Wahyu Setyawan, Shidong Wang, Junkai Xue, Kesong Yang, Richard H Taylor, Lance J Nelson, Gus L W Hart, Stefano Sanvito, Marco Buongiorno-Nardelli, Natalio Mingo, and Ohad Levy. AFLOWLIB.ORG: A distributed materials properties repository from high-throughput ab initio calculations. *Computational Materials Science*, 58:227–235, 2012. doi: 10.1016/j.commatsci.2012.02.002.

[16] Daniel W Davies, Keith T Butler, Adam J Jackson, Andrew Morris, Jarvist M Frost, Jonathan M Skelton, and Aron Walsh. Computational Screening of All Stoichiometric Inorganic Materials. *Chem*, 1(4):617–627, 2016. ISSN 2451-9294. doi: <https://doi.org/10.1016/j.chempr.2016.09.010>. URL <https://www.sciencedirect.com/science/article/pii/S2451929416301553>.

[17] Bowen Deng, Yunyeong Choi, Peichen Zhong, Janosh Riebesell, Shashwat Anand, Zhuohan Li, KyuJung Jun, Kristin A. Persson, and Gerbrand Ceder. Systematic softening in universal machine learning interatomic potentials. *npj Computational Materials*, 11(1):9, 2025. doi: 10.1038/s41524-024-01500-6. URL <https://doi.org/10.1038/s41524-024-01500-6>.

[18] Max Welling Diederik P. Kingma. Auto-encoding variational bayes. In *International Conference on Learning Representations*, 2014.

[19] Hong Ding, Shyam S Dwaraknath, Lauren Garten, Paul Ndione, David Ginley, and Kristin A Persson. Computational Approach for Epitaxial Polymorph Stabilization through Substrate Selection. *ACS Applied Materials & Interfaces*, 8(20):13086–13093, 2016. doi: 10.1021/acsami.6b01630. URL <https://doi.org/10.1021/acsami.6b01630>.

[20] Alexandre Duval, Victor Schmidt, Alex Hernandez Garcia, Santiago Miret, Fragkiskos D Malliaros, Yoshua Bengio, and David Rolnick. FAENet: Frame Averaging Equivariant GNN for Materials Modeling. In *Proceedings of the 40th International Conference on Machine Learning*, pages 9013–9033, 2023.

[21] Matthias Fey and Jan E. Lenssen. Fast graph representation learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019.

[22] Scott Fredericks, Kevin Parrish, Dean Sayre, and Qiang Zhu. PyXtal: A Python library for crystal structure generation and symmetry analysis. *Computer Physics Communications*, 261:107810, 2021. ISSN 0010-4655. doi: <https://doi.org/10.1016/j.cpc.2020.107810>. URL <https://www.sciencedirect.com/science/article/pii/S0010465520304057>.

[23] Jingheng Fu, Mikael Kuisma, Ask Hjorth Larsen, Kohei Shinohara, Atsushi Togo, and Kristian S Thygesen. Symmetry classification of 2d materials: layer groups versus space groups. *2D Materials*, 11:035009, 2024. doi: 10.1088/2053-1583/ad3e0c.

[24] Xiang Fu, Brandon M. Wood, Luis Barroso-Luque, Daniel S. Levine, Meng Gao, Misko Dzamba, and C. Lawrence Zitnick. Learning smooth and expressive interatomic potentials for physical property prediction, 2025. URL <https://arxiv.org/abs/2502.12147>.

[25] Max C. Gallant, Matthew J. McDermott, Bryant Li, and Kristin A. Persson. A cellular automaton simulation for predicting phase evolution in solid-state reactions. *Chemistry of Materials*, 37(1):210–223, 2025. doi: 10.1021/acs.chemmater.4c02301. URL <https://doi.org/10.1021/acs.chemmater.4c02301>.

[26] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative Adversarial Nets. In Z Ghahramani, M Welling, C Cortes, N Lawrence, and K Q Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014. URL https://proceedings.neurips.cc/paper{_}files/paper/2014/file/5ca3e9b122f61f8f06494c97b1afccf3-Paper.pdf.

[27] Alex Graves, Rupesh Kumar Srivastava, Timothy Atkinson, and Faustino J. Gomez. Bayesian flow networks. *CorR*, abs/2308.07037, 2023. URL <https://doi.org/10.48550/arXiv.2308.07037>.

[28] Martin A. Green, Anita Ho-Baillie, and Henry J. Snaith. The emergence of perovskite solar cells. *Nature Photonics* 2014 8:7, 8(7):506–514, 2014. ISSN 1749-4893. doi: 10.1038/nphoton.2014.134. URL <https://www.nature.com/articles/nphoton.2014.134>.

[29] R. W. Grosse-Kunstleve, N. K. Sauter, and P. D. Adams. Numerically stable algorithms for the computation of reduced unit cells. *Acta Cryst.*, 60(1):1–6, 2004. ISSN 0108-7673. doi: 10.1107/S010876730302186X. URL <https://journals.iucr.org/paper?sh5006><https://journals.iucr.org/a/issues/2004/01/00/sh5006/>.

[30] Ralf W. Grosse-Kunstleve, Buddy Wong, Marat Mustyakimov, and Paul D. Adams. Exact direct-space asymmetric units for the 230 crystallographic space groups. *Acta Crystallogr A*, 67(3):269–275, 2011. ISSN 0108-7673. doi: 10.1107/S0108767311007008. URL <https://journals.iucr.org/paper?pz5088>.

[31] Nate Gruver, Anuroop Sriram, Andrea Madotto, Andrew Gordon Wilson, C. Lawrence Zitnick, and Zachary Ward Ulissi. Fine-tuned language models generate stable inorganic materials as text. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=vN9fpfqoP1>.

[32] Yu Han, Chi Ding, Junjie Wang, Hao Gao, Jiuyang Shi, Shaobo Yu, Qiuhan Jia, Shuning Pan, and Jian Sun. Efficient crystal structure prediction based on the symmetry principle. *Nature Computational Science*, 5:255–267, 2025. ISSN 2662-8457. doi: 10.1038/s43588-025-00775-z. URL <https://www.nature.com/articles/s43588-025-00775-z>.

[33] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising Diffusion Probabilistic Models. In H Larochelle, M Ranzato, R Hadsell, M F Balcan, and H Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 6840–6851. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper{_}files/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf.

[34] Jordan Hoffmann, Louis Maestrati, Yoshihide Sawada, Jian Tang, Jean Michel Sellier, and Yoshua Bengio. Data-driven approach to encoding and decoding 3-D crystal structures, 2019. URL <https://arxiv.org/abs/1909.00949>.

[35] Philipp Höllmer, Thomas Egg, Maya Martirosyan, Eric Fuemmeler, Amit Gupta, Zeren Shui, Pawan Prakash, Adrian Roitberg, Mingjie Liu, George Karypis, Mark Transtrum, Richard Hennig, Ellad B. Tadmor, and Stefano Martiniani. Open materials generation with stochastic interpolants. In *AI for Accelerated Materials Design - ICLR 2025*, 2025. URL <https://openreview.net/forum?id=ka2jxtCrkK>.

[36] Aapo Hyvärinen. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(24):695–709, 2005. URL <http://jmlr.org/papers/v6/hyvarinen05a.html>.

[37] Anubhav Jain, Geoffroy Hautier, Charles J. Moore, Shyue Ping Ong, Christopher C. Fischer, Tim Mueller, Kristin A. Persson, and Gerbrand Ceder. A high-throughput infrastructure for density functional theory calculations. *Computational Materials Science*, 50(8):2295–2310, 2011. ISSN 0927-0256. doi: <https://doi.org/10.1016/j.commatsci.2011.02.023>. URL <https://www.sciencedirect.com/science/article/pii/S0927025611001133>.

[38] Anubhav Jain, Geoffroy Hautier, Shyue Ping Ong, Charles J. Moore, Christopher C. Fischer, Kristin A. Persson, and Gerbrand Ceder. Formation enthalpies by mixing gga and gga + u calculations. *Phys. Rev. B*, 84:045115, Jul 2011. doi: 10.1103/PhysRevB.84.045115. URL <https://link.aps.org/doi/10.1103/PhysRevB.84.045115>.

[39] Anubhav Jain, Ping Ong, Geoffroy Hautier, Wei Chen, William Davidson Richards, Stephen Dacek, Shreyas Cholia, Dan Gunter, David Skinner, Gerbrand Ceder, and Kristin A Persson. Commentary: The Materials Project: A materials genome approach to accelerating materials innovation. *APL Mater*, 1:11002, 2013. doi: 10.1063/1.4812323. URL <https://doi.org/10.1063/1.4812323>.

[40] Rui Jiao, Wenbing Huang, Peijia Lin, Jiaqi Han, Pin Chen, Yutong Lu, and Yang Liu. Crystal Structure Prediction by Joint Equivariant Diffusion. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.

[41] Rui Jiao, Wenbing Huang, Yu Liu, Deli Zhao, and Yang Liu. Space group constrained crystal generation. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=jkvZ7v40mP>.

[42] Chaitanya K. Joshi, Xiang Fu, Yi-Lun Liao, Vahe Gharakhanyan, Benjamin Kurt Miller, Anuroop Sriram, and Zachary Ward Ulissi. All-atom diffusion transformers. In *AI for Accelerated Materials Design - ICLR 2025*, 2025. URL <https://openreview.net/forum?id=mXApCXR21F>.

[43] Vir Karan, Max C. Gallant, Yuxing Fei, Gerbrand Ceder, and Kristin A. Persson. Ion correlations explain kinetic selectivity in diffusion-limited solid state synthesis reactions, 2025. URL <https://arxiv.org/abs/2501.08560>.

[44] Isay Katsman, Aaron Lou, Derek Lim, Qingxuan Jiang, Ser Nam Lim, and Christopher M De Sa. Equivariant manifold flows. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 10600–10612. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/581b41df0cd50ace849e061ef74827fc-Paper.pdf.

[45] Nikita Kazeev, Ruiming Zhu, Ignat Romanov, Andrey E Ustyuzhanin, Shuya Yamazaki, Wei Nong, and Kedar Hippalgaonkar. Wyckofftransformer: Generation of symmetric crystals. In *AI for Accelerated Materials Design - NeurIPS 2024*, 2024. URL <https://openreview.net/forum?id=Jcy1bP0qrY>.

[46] Nikita Kazeev, Ruiming Zhu, Ignat Romanov, Andrey E Ustyuzhanin, Shuya Yamazaki, Wei Nong, and Kedar Hippalgaonkar. Wyckofftransformer: Generation of symmetric crystals. In *AI for Accelerated Materials Design - NeurIPS 2024*, 2024. URL <https://openreview.net/forum?id=Jcy1bP0qrY>.

[47] Filip Ekström Kelvinius, Oskar B. Andersson, Abhijith S. Parackal, Dong Qian, Rickard Armiento, and Fredrik Lindsten. Wyckoffdiff – a generative diffusion model for crystal symmetry, 2025. URL <https://arxiv.org/abs/2502.06485>.

[48] Joonsoo Kim, Yuki Kimura, Brian Puchala, Tomoya Yamazaki, Udo Becker, and Wenhao Sun. Dissolution enables dolomite crystal growth near ambient conditions. *Science*, 382(6673):915–920, 2023. doi: 10.1126/science.adi3690. URL <https://www.science.org/doi/abs/10.1126/science.adi3690>.

[49] Sungwon Kim, Juhwan Noh, Geun Ho Gu, Alán Aspuru-Guzik, and Yousung Jung. Generative Adversarial Networks for Crystal Structure Prediction. *ACS Central Science*, 6(8):1412–1420, 2020. URL <http://arxiv.org/abs/2004.01396>.

[50] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR (Poster)*, 2015. URL <http://arxiv.org/abs/1412.6980>.

[51] Johannes Klicpera, Florian Becker, and Stephan Günnemann. Gemnet: Universal directional graph neural networks for molecules. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021. URL https://openreview.net/forum?id=HS_s0axS9K-.

[52] Jonas Köhler, Leon Klein, and Frank Noe. Equivariant flows: Exact likelihood generative learning for symmetric densities. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 5361–5370. PMLR, 13–18 Jul 2020. URL <https://proceedings.mlr.press/v119/kohler20a.html>.

[53] Tobias Kreiman and Aditi S. Krishnapriyan. Understanding and mitigating distribution shifts for machine learning force fields, 2025. URL <https://arxiv.org/abs/2503.08674>.

[54] Georg Kresse and Jürgen Furthmüller. Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set. *Physical review B*, 54(16):11169, 1996.

[55] Daniel Levy, Siba Smarak Panigrahi, Sékou-Oumar Kaba, Qiang Zhu, Kin Long Kelvin Lee, Mikhail Galkin, Santiago Miret, and Siamak Ravanbakhsh. SymmCD: Symmetry-preserving crystal generation with diffusion models. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=xnssGv9rpW>.

[56] Qi Li, Rui Jiao, Liming Wu, Tiannian Zhu, Wenbing Huang, Shifeng Jin, Yang Liu, Hongming Weng, and Xiaolong Chen. Powder diffraction crystal structure determination using generative models, 2024. URL <https://arxiv.org/abs/2409.04727>.

[57] Peijia Lin, Pin Chen, Rui Jiao, Qing Mo, Cen Jianhuan, Wenbing Huang, Yang Liu, Dan Huang, and Yutong Lu. Equivariant diffusion for crystal structure prediction. In *Forty-first International Conference on Machine Learning*, 2024. URL <https://openreview.net/forum?id=VRv8KjJNuJ>.

[58] Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=PqvMRDCJT9t>.

[59] Yuchen Lou and Alex M. Ganose. Discovery of highly anisotropic dielectric crystals with equivariant graph neural networks. *Faraday Discussions*, 256:255–274, 2025. doi: 10.1039/D4FD00096J.

[60] Matthew J McDermott, Brennan C McBride, Corlyn E Regier, Gia Thinh Tran, Yu Chen, Adam A Corrao, Max C Gallant, Gabrielle E Kamm, Christopher J Bartel, Karena W Chapman, Peter G Khalifah, Gerbrand Ceder, James R Neilson, and Kristin A Persson. Assessing Thermodynamic Selectivity of Solid-State Reactions for the Predictive Synthesis of Inorganic Materials. *ACS Central Science*, 9(10):1957–1975, 2023. ISSN 2374-7943. doi: 10.1021/acscentsci.3c01051. URL <https://doi.org/10.1021/acscentsci.3c01051>.

[61] Aaron Meurer, Christopher P. Smith, Mateusz Paprocki, Ondřej Čertík, Sergey B. Kirpichev, Matthew Rocklin, AMiT Kumar, Sergiu Ivanov, Jason K. Moore, Sartaj Singh, Thilina Rathnayake, Sean Vig, Brian E. Granger, Richard P. Muller, Francesco Bonazzi, Harsh Gupta, Shivam Vats, Fredrik Johansson, Fabian Pedregosa, Matthew J. Curry, Andy R. Terrel, Štěpán Roučka, Ashutosh Saboo, Isuru Fernando, Sumith Kulal, Robert Cimrman, and Anthony Scopatz. Sympy: symbolic computing in python. *PeerJ Computer Science*, 3:e103, 2017. ISSN 2376-5992. doi: 10.7717/peerj-cs.103. URL <https://doi.org/10.7717/peerj-cs.103>.

[62] Benjamin Kurt Miller, Ricky T.Q. Chen, Anuroop Sriram, and Brandon M. Wood. FlowMM: Generating Materials with Riemannian Flow Matching. *Proceedings of Machine Learning Research*, 235:35664–35686, 2024. ISSN 26403498. URL <https://arxiv.org/abs/2406.04713v1>.

[63] Mehran Mirramezani, Anne S. Meeussen, Katia Bertoldi, Peter Orbantz, and Ryan P. Adams. Designing mechanical meta-materials by learning equivariant flows, 2024. URL <https://arxiv.org/abs/2410.02385>.

[64] Mistal, Alex Hernández-García, Alexandra Volokhova, Alexandre AGM Duval, Yoshua Bengio, Divya Sharma, Pierre Luc Carrier, Michał Koziarski, and Victor Schmidt. Crystal-GFN: sampling materials with desirable properties and constraints. In *AI for Accelerated Materials Design - NeurIPS 2023 Workshop*, 2023. URL <https://openreview.net/forum?id=1167FjdPOv>.

[65] F.E. Neumann. *Vorlesungen über die Theorie der Elasticität der festen Körper und des Lichtäthers, gehalten an der Universität Königsberg*. Leipzig, G. G., 1885.

[66] Naoki Nitta, Feixiang Wu, Jung Tae Lee, and Gleb Yushin. Li-ion battery materials: present and future. *Materials Today*, 18(5):252–264, 2015. ISSN 1369-7021. doi: 10.1016/J.MATTOD.2014.10.040.

[67] Juhwan Noh, Jaehoon Kim, Helge S Stein, Benjamin Sanchez-Lengeling, John M Gregoire, Alan Aspuru-Guzik, and Yousung Jung. Inverse Design of Solid-State Materials via a Continuous Representation. *Matter*, 1:1370–1384, 2019. doi: 10.1016/j.matt.2019.08.017. URL <https://doi.org/10.1016/j.matt.2019.08.017>.

[68] Andrew Novick, Quan Nguyen, Matthew Jankousky, M. Brooks Tellekamp, Eric S. Toberer, and Vladan Stevanović. Basin-size mapping: Prediction of metastable polymorph synthesizability across tac-tan alloys. *Journal of the American Chemical Society*, 147(5):4419–4429, 2025. doi: 10.1021/jacs.4c15441. URL <https://doi.org/10.1021/jacs.4c15441>. PMID: 39835391.

[69] Shyue Ping Ong, William Davidson Richards, Anubhav Jain, Geoffroy Hautier, Michael Kocher, Shreyas Cholia, Dan Gunter, Vincent L Chevrier, Kristin A Persson, and Gerbrand Ceder. Python materials genomics (pymatgen): A robust, open-source python library for materials analysis. *Computational Materials Science*, 68:314–319, 2013.

[70] Shyue Ping Ong, William Davidson Richards, Anubhav Jain, Geoffroy Hautier, Michael Kocher, Shreyas Cholia, Dan Gunter, Vincent L. Chevrier, Kristin A. Persson, and Gerbrand Ceder. Python Materials Genomics (pymatgen): A robust, open-source python library for materials analysis. *Computational Materials Science*, 68:314–319, 2013. ISSN 09270256. doi: 10.1016/j.commatsci.2012.10.028.

[71] Hyunsoo Park, Anthony Onwuli, and Aron Walsh. Exploration of crystal chemical space using text-guided generative artificial intelligence. *Nature Communications*, 16(1):4379, 2025. doi: 10.1038/s41467-025-59636-y. URL <https://doi.org/10.1038/s41467-025-59636-y>.

[72] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *NIPS 2017 Workshop Autodiff*, 2017.

[73] John P Perdew, Kieron Burke, and Matthias Ernzerhof. Generalized gradient approximation made simple. *Physical review letters*, 77(18):3865, 1996.

[74] Omri Puny, Matan Atzmon, Edward J. Smith, Ishan Misra, Aditya Grover, Heli Ben-Hamu, and Yaron Lipman. Frame averaging for invariant and equivariant network design. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=zIUyj55nXR>.

[75] Zekun Ren, Juhwan Noh, Siyu Tian, Felipe Oviedo, Guangzong Xing, Qiaohao Liang, Armin Aberle, Yi Liu, Qianxiao Li, Senthilnath Jayavelu, Kedar Hippalgaonkar, Yousung Jung, and Tonio Buonassisi. Inverse design of crystals using generalized invertible crystallographic representation. *Matter*, 5(1), 2022. URL <https://arxiv.org/abs/2005.07609>.

[76] Danilo Jimenez Rezende, Sébastien Racanière, Irina Higgins, and Peter Toth. Equivariant hamiltonian flows, 2019. URL <https://arxiv.org/abs/1909.13739>.

[77] Benjamin Rhodes, Sander Vandenhaut, Vaidotas Šimkus, James Gin, Jonathan Godwin, Tim Dugnan, and Mark Neumann. Orb-v3: atomistic simulation at scale, 2025. URL <https://arxiv.org/abs/2504.06231>.

[78] Laura Ruple, Luca Torresi, Henrik Schopmans, and Pascal Friederich. Symmetry-aware bayesian flow networks for crystal generation, 2025. URL <https://arxiv.org/abs/2502.03146>.

[79] James E Saal, Scott Kirklin, Muratahan Aykol, Bryce Meredig, and C Wolverton. Materials Design and Discovery with High-Throughput Density Functional Theory: The Open Quantum Materials Database (OQMD). *JOM*, 65:1501–1509, 2013. doi: 10.1007/s11837-013-0755-4.

[80] Victor Garcia Satorras, Emiel Hoogeboom, and Max Welling. E(n) equivariant graph neural networks. In *ICML*, pages 9323–9332, 2021. URL <http://proceedings.mlr.press/v139/satorras21a.html>.

[81] Lisa Schneckenreiter, Richard Freinschlag, Florian Sestak, Johannes Brandstetter, Günter Klambauer, and Andreas Mayr. GNN-VPA: A variance-preserving aggregation strategy for graph neural networks. In *5th Workshop on practical ML for limited/low resource settings*, 2024. URL <https://openreview.net/forum?id=cmrRmu2afD>.

[82] Sharon Shwartz, Raoul Weil, Mordechai Segev, Eugene Lakin, Emil Zolotoyabko, Vinod M Menon, Stephen R Forrest, and Uri El-Hanany. Light-induced symmetry breaking and related giant enhancement of nonlinear properties in CdZnTe:V crystals. *Opt. Express*, 14(20):9385–9390, 2006. doi: 10.1364/OE.14.009385. URL <https://opg.optica.org/oe/abstract.cfm?URI=oe-14-20-9385>.

[83] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 2256–2265, Lille, France, 2015. PMLR. URL <https://proceedings.mlr.press/v37/sohl-dickstein15.html>.

[84] Yang Song and Stefano Ermon. Generative Modeling by Estimating Gradients of the Data Distribution. In *Advances in Neural Information Processing Systems*, 2019.

[85] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=PxTIG12RRHS>.

[86] Anuroop Sriram, Benjamin Kurt Miller, Ricky T. Q. Chen, and Brandon M Wood. FlowLLM: Flow matching for material generation with large language models as base distributions. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL <https://openreview.net/forum?id=0bFXbEMz8e>.

[87] Vladan Stevanović, Stephan Lany, Xiuwen Zhang, and Alex Zunger. Correcting density functional theory for accurate predictions of compound enthalpies of formation: Fitted elemental-phase reference energies. *Phys. Rev. B*, 85:115104, Mar 2012. doi: 10.1103/PhysRevB.85.115104. URL <https://link.aps.org/doi/10.1103/PhysRevB.85.115104>.

[88] Wenhao Sun, Stephen T. Dacek, Shyue Ping Ong, Geoffroy Hautier, Anubhav Jain, William D. Richards, Anthony C. Gamst, Kristin A. Persson, and Gerbrand Ceder. The thermodynamic scale of inorganic crystalline metastability. *Science Advances*, 2(11):e1600225, 2016. doi: 10.1126/sciadv.1600225. URL <https://www.science.org/doi/10.1126/sciadv.1600225>.

[89] Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier Features Let Networks Learn High Frequency Functions in Low Dimensional Domains. In H Larochelle, M Ranzato, R Hadsell, M F Balcan, and H Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 7537–7547. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper{_}files/paper/2020/file/55053683268957697aa39fba6f231c68-Paper.pdf.

[90] Feng Tang, Hoi Chun Po, Ashvin Vishwanath, and Xiangang Wan. Comprehensive search for topological materials using symmetry indicators. *Nature*, 566(7745):486–489, 2019. ISSN 1476-4687. doi: 10.1038/s41586-019-0937-5. URL <https://www.nature.com/articles/s41586-019-0937-5>.

[91] Atsushi Togo, Kohei Shinohara, and Isao Tanaka. Spglib: a software library for crystal symmetry search. *Science and Technology of Advanced Materials: Methods*, 4:1, 2024. doi: 10.1080/27660400.2024.2384822. URL <https://www.tandfonline.com/doi/abs/10.1080/27660400.2024.2384822>.

[92] Ashish Vaswani, Google Brain, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention Is All You Need. In *NIPS*, 2017.

[93] Logan Ward, Ankit Agrawal, Alok Choudhary, and Christopher Wolverton. A general-purpose machine learning framework for predicting properties of inorganic materials. *npj Computational Materials*, 2(1):16028, 2016. ISSN 2057-3960. doi: 10.1038/npjcompumats.2016.28. URL <https://doi.org/10.1038/npjcompumats.2016.28>.

[94] Haruki Watanabe, Hoi Chun Po, and Ashvin Vishwanath. Structure and topology of band structures in the 1651 magnetic space groups. *Science Advances*, 4(8):eaat8685, 2018. doi: 10.1126/sciadv.aat8685. URL <https://www.science.org/doi/abs/10.1126/sciadv.aat8685>.

[95] Mingjian Wen, Matthew K. Horton, Jason M. Munro, Patrick Huck, and Kristin A. Persson. An equivariant graph neural network for the elasticity tensors of all seven crystal systems. *Digital Discovery*, 3:869–882, 2024. doi: 10.1039/D3DD00233K. URL <http://dx.doi.org/10.1039/D3DD00233K>.

[96] Rachel Woods-Robinson, Yanbing Han, Hanyu Zhang, Tursun Ablekim, Imran Khan, Kristin A. Persson, and Andriy Zakutayev. Wide Band Gap Chalcogenide Semiconductors. *Chemical Reviews*, 120(9):4007–4055, 2020. ISSN 15206890. doi: 10.1021/ACS.CHEMREV.9B00600. URL <https://pubs.acs.org/doi/full/10.1021/acs.chemrev.9b00600>.

[97] Hanlin Wu, Yuxuan Song, Jingjing Gong, Ziyao Cao, Yawen Ouyang, Jianbing Zhang, Hao Zhou, Wei-Ying Ma, and Jingjing Liu. A periodic bayesian flow for material generation. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=Lz0XW99tE0>.

[98] Zhenyu Xiao, Jianzhou Zhao, Yanqi Li, Ryuichi Shindou, and Zhi-Da Song. Spin space groups: Full classification and applications. *Phys. Rev. X*, 14:031037, Aug 2024. doi: 10.1103/PhysRevX.14.031037. URL <https://link.aps.org/doi/10.1103/PhysRevX.14.031037>.

[99] Tian Xie and Jeffrey C. Grossman. Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties. *Phys. Rev. Lett.*, 120:145301, 2018. doi: 10.1103/PhysRevLett.120.145301. URL <https://link.aps.org/doi/10.1103/PhysRevLett.120.145301>.

[100] Tian Xie, Xiang Fu, Octavian-Eugen Ganea, Regina Barzilay, and Tommi Jaakkola. Crystal Diffusion Variational Autoencoder for Periodic Material Generation. In *International Conference on Learning Representations*, 2022.

[101] Keqiang Yan, Xiner Li, Hongyi Ling, Kenna Ashen, Carl Edwards, Raymundo Arroyave, Marinka Zitnik, Heng Ji, Xiaofeng Qian, Xiaoning Qian, and Shuiwang Ji. Invariant tokenization of crystalline materials for language model enabled generation. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL <https://openreview.net/forum?id=18FGRNd0wZ>.

[102] Keqiang Yan, Alexandra Saxton, Xiaofeng Qian, Xiaoning Qian, and Shuiwang Ji. A space group symmetry informed network for $o(3)$ equivariant crystal tensor prediction. In *Forty-first International Conference on Machine Learning*, 2024. URL <https://openreview.net/forum?id=BOFjRnJ9mX>.

- [103] Sherry Yang, Simon Batzner, Ruiqi Gao, Muratahan Aykol, Alexander L Gaunt, Brendan McMorrow, Danilo J Rezende, Dale Schuurmans, Igor Mordatch, and Ekin D Cubuk. Generative hierarchical materials search. *arXiv preprint arXiv:2409.06762*, 2024.
- [104] Sherry Yang, KwangHwan Cho, Amil Merchant, Pieter Abbeel, Dale Schuurmans, Igor Mordatch, and Ekin Dogus Cubuk. Scalable diffusion for materials generation. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=wm4W1HoXpC>.
- [105] Dmitry Yarotsky. Universal Approximations of Invariant Maps by Neural Networks. *Constructive Approximation*, 55(1):407–474, 2022. ISSN 1432-0940. doi: 10.1007/s00365-021-09546-1. URL <https://doi.org/10.1007/s00365-021-09546-1>.
- [106] Werner Zellinger, Edwin Lughofe, Susanne Saminger-Platz, Thomas Grubinger, and Thomas Natschläger. Central Moment Discrepancy (CMD) for Domain-Invariant Representation Learning. In *International Conference on Learning Representations*, 2017. URL <http://www.scch.at>.
- [107] Claudio Zeni, Robert Pinsler, Daniel Zügner, Andrew Fowler, Matthew Horton, Xiang Fu, Aliaksandra Shysheya, Jonathan Crabbé, Lixin Sun, Jake Smith, Bichlien Nguyen, Hannes Schulz, Sarah Lewis, Chin-Wei Huang, Ziheng Lu, Yichi Zhou, Han Yang, Hongxia Hao, Jielan Li, Ryota Tomioka, and Tian Xie. A generative model for inorganic materials design. *Nature*, 2025. URL <https://arxiv.org/abs/2312.03687v2>.
- [108] Ruiming Zhu, Wei Nong, Shuya Yamazaki, and Kedar Hippalgaonkar. WyCryst: Wyckoff inorganic crystal generator framework. *Matter*, 7(10):3469–3488, 2024. ISSN 2590-2393. doi: 10.1016/j.matt.2024.05.042. URL <https://doi.org/10.1016/j.matt.2024.05.042>.
- [109] Nils E.R. Zimmermann and Anubhav Jain. Local structure order parameters and site fingerprints for quantification of coordination environment and crystal structure similarity. *RSC Advances*, 10(10):6063–6081, 2020. ISSN 20462069. doi: 10.1039/c9ra07755c. URL <https://pubs.rsc.org/en/content/articlehtml/2020/ra/c9ra07755c><https://pubs.rsc.org/en/content/articlelanding/2020/ra/c9ra07755c>.
- [110] A. Zur and T. C. McGill. Lattice match: An application to heteroepitaxy. *Journal of Applied Physics*, 55(2):378–386, 01 1984. ISSN 0021-8979. doi: 10.1063/1.333084. URL <https://doi.org/10.1063/1.333084>.

A Appendix

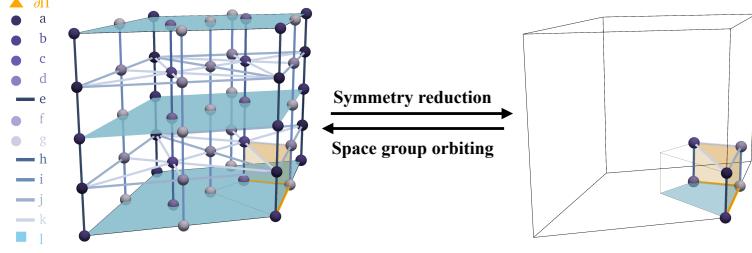


Figure 3: Special Wyckoff positions and the asymmetric unit in the conventional unit cell of hexagonal space group 192. Closed asymmetric unit boundary edges and facets ($\partial\Pi$) are in orange.

Table 4: Normalized root mean square displacements (RMSD) in Angstroms and their standard deviations calculated with pymatgen StructureMatcher between as-generated and DFT-relaxed crystals. RMSD is normalized by $(V/N)^{1/3}$ for unit cell volume V and number of atoms per cell N .

CDVAE [100]	DiffCSP++ [41]	MatterGen [107]	SGEquiDiff	SGEquiDiff+LDiff
0.0642 ± 0.0568	0.0543 ± 0.0812	0.0270 ± 0.0387	0.0619 ± 0.0804	0.0860 ± 0.0810

A.1 Proofs

Theorem 1. Let G be a space group, $x \in \mathbb{R}^3$ be a point residing in a Wyckoff position with stabilizer group G_x , and $f : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ be a space group equivariant function. Then, for arbitrary constant $c \in \mathbb{R}$ and space group element $g \in G_x$,

$$g(x + cf(x)) = x + cf(x), \quad (9)$$

or equivalently, $G_x \subseteq G_{x+cf(x)}$.

Proof. Let $\{R_{g_x}, v_{g_x}\} = g_x \in G_x$ be a stabilizer group element where $R_{g_x} \in O(3)$ is a point group operation, $v_{g_x} \in [0, 1]^3$ is a fractional lattice translation, and the action of g_x is $R_{g_x}(\cdot) + v_{g_x}$. The definition of a stabilizer group requires that

$$R_{g_x}x + v_{g_x} = x \implies v_{g_x} = 0.$$

Thus, we have

$$f(R_{g_x}x) = f(x). \quad (10)$$

The definition of space group equivariance requires that for any $R \in G$,

$$f(Rx) = Rf(x). \quad (11)$$

Combining equations 10 and 11, we see

$$R_{g_x}f(x) = f(x), \quad (12)$$

i.e., $f(x)$ is also stabilized by G_x . Finally, we have

$$R_{g_x}(x + cf(x)) = R_{g_x}x + cR_{g_x}f(x) \quad (13)$$

$$= x + cf(x). \quad (14)$$

Theorem 2. The score of the SGWN distribution is space group equivariant, i.e., given $g = \{R, v\} \in G$,

$$\nabla_{x_t} \log p(Rx_t + v | x_0) = R \nabla_{x_t} \log p(x_t | x_0).$$

Proof. The score of the SGWN is given as

$$\nabla_{x_t} \log p(x_t|x_0) = \frac{\sum_{g_j=\{R_j, t_j\} \in G} \exp\left(\frac{-\|x_t - (R_j x_0 + t_j)\|_F^2}{2\sigma_t^2}\right) (R_j x_0 + t_j - x_t)}{\sigma_t^2 p(x_t|x_0)}.$$

Now, we transform the input with $g = \{R, v\}$.

$$\nabla_{x_t} \log p(Rx_t + v|x_0) = \frac{\sum_{g_j=\{R_j, t_j\} \in G} \exp\left(\frac{-\|Rx_t + v - (R_j x_0 + t_j)\|_F^2}{2\sigma_t^2}\right) (R_j x_0 + t_j - Rx_t - v)}{\sigma_t^2 p(Rx_t + v|x_0)}.$$

We will prove the equivariance of the numerator, which we denote as $\phi(Rx_t + v|x_0)$. For the denominator, it can be similarly proved that the SGWN is G -invariant, i.e., $p(Rx_t + v|x_0) = p(x_t|x_0)$.

$$\begin{aligned} \phi(Rx_t + v|x_0) &= \sum_{\{R_j, t_j\}} \exp\left(\frac{-\|Rx_t + v - (R_j x_0 + t_j)\|_F^2}{2\sigma_t^2}\right) (R_j x_0 + t_j - Rx_t - v) \\ &= \sum_{\{R_j, t_j\}} \exp\left(\frac{-\|R(x_t - R^T(R_j x_0 + t_j - v))\|_F^2}{2\sigma_t^2}\right) [R(R^T(R_j x_0 + t_j - v) - x_t)] \\ &= \sum_{\{R_j, t_j\}} \exp\left(\frac{-\|x_t - R^T(R_j x_0 + t_j - v)\|_F^2}{2\sigma_t^2}\right) [R(R^T(R_j x_0 + t_j - v) - x_t)] \\ &= \sum_{\{R_j, t_j\}} \exp\left(\frac{-\|x_t - g^{-1}(R_j x_0 + t_j)\|_F^2}{2\sigma_t^2}\right) [R(g^{-1}(R_j x_0 + t_j) - x_t)] \\ &= \sum_{\{R_{j'}, t_{j'}\}} \exp\left(\frac{-\|x_t - (R_{j'} x_0 + t_{j'})\|_F^2}{2\sigma_t^2}\right) [R(R_{j'} x_0 + t_{j'} - x_t)] \\ &= R\phi(x_t|x_0). \end{aligned}$$

A.2 Implementation details

We approximated the SGWN distribution with a truncated sum over lattice translations as

$$q(x_t|x_0) \propto \sum_{t_L \in \mathbb{Z}^3 \cap [-m, m]^3} \sum_{\{R_j, t_j\} \in G/T_L} \exp\left(\frac{-\|x_t - (R_j x_0 + t_j + t_L)\|_F^2}{2\sigma_t^2}\right)$$

where $m \in \mathbb{Z}_+$. We have dropped the dependence of q on G and σ_t from the notation for clarity. The space group-equivariant score of $q(x_t|x_0)$ is given as

$$\begin{aligned} \nabla_{x_t} \log q(x_t|x_0) &= \nabla_{x_t} \log \left[\sum_{t_L \in \mathbb{Z}^3 \cap [-m, m]^3} \sum_{\{R_j, t_j\}} \exp\left(\frac{-\|x_t - (R_j x_0 + t_j + t_L)\|_F^2}{2\sigma_t^2}\right) \right] \\ &= \frac{\sum_{t_L \in \mathbb{Z}^3 \cap [-m, m]^3} \sum_{\{R_j, t_j\}} \exp\left(\frac{-\|x_t - (R_j x_0 + t_j + t_L)\|_F^2}{2\sigma_t^2}\right) (R_j x_0 + t_j + t_L - x_t)}{\sigma_t^2 q(x_t|x_0)}. \end{aligned} \tag{15}$$

For each Wyckoff position w , we pre-computed approximate values of $\lambda_t = \mathbb{E}_{x_t \sim p(x_t|x_0), x_0 \sim \mathcal{P}_w(\text{Uniform}(\cdot))} [\|\nabla_{x_t} \log p(x_t|x_0)\|]^{-1}$ with Monte Carlo sampling. We denote $\mathcal{P}_w(\text{Uniform}(\cdot))$ as the uniform distribution on the Wyckoff shapes (see Sec. 4.1 and A.3) corresponding to Wyckoff position w . First, we sampled s points $\{x_0^i\}_{i=1}^s$ from $\mathcal{P}_w(\text{Uniform}(\cdot))$. Then, for each sample x_0^i , we sampled Gaussian noise $\epsilon \sim \mathcal{N}(0, \sigma_t^2)$ and reparameterized as $x_t^i = x_0^i + \mathcal{P}_{x_0^i}(\epsilon)$. The approximation of λ_t for Wyckoff position w in space group G was then computed as

$$\tilde{\lambda}_t = \left[\frac{1}{s} \sum_{i=1}^s \nabla_{x_t^i} \log q(x_t^i|x_0^i; G, w, \sigma_t) \right]^{-1}$$

In our experiments, we set $s = 2500$, $m = 3$, $\sigma_1 = 0.002$, $\sigma_T = 0.5$, and $T = 1000$.

A.3 Uniform priors on Wyckoff positions

To sample from the uniform distribution in a given Wyckoff position, we sampled uniformly on the Wyckoff shapes in the ASU. The likelihoods of these uniform priors are simply given as 1 for 0D Wyckoff positions, and the reciprocal of the total length, area, and volume of the relevant Wyckoff shapes for the 1-, 2-, and 3-D Wyckoff positions, respectively. For a 0D Wyckoff position, coordinates are fully determined and ‘sampling’ simply returns a point. For a 1D Wyckoff position, we sampled a line segment with probability proportional to its length, and then sampled uniformly on the line segment. For a 2D Wyckoff position, we pre-computed Delaunay triangulations of the polygonal facets belonging to the Wyckoff position, sampled a triangle proportionally to its area, and then sampled from a uniform Dirichlet distribution in the triangle. For a 3D Wyckoff position, we used rejection sampling from a bounding box around the ASU.

A.4 Architecture

A.4.1 Wyckoff-Element Transformer

Our encoder-decoder Transformer architecture can be summarized as follows:

$$\begin{aligned}
z^0 &\leftarrow \text{MLP}(e_{SG} || e_W || e_A) \\
z^{l+1} &\leftarrow \text{EncoderLayer}(z_i^l, m_{\text{causal}}) \\
z_W, z_A &\leftarrow \text{Split}(z^{l_{\text{max}}}) \\
z_A &\leftarrow \text{MLP}(z_A, e_W) \\
p_{W,\text{stop}} &\leftarrow \text{Attention}(K = [e_W^{\text{all}} || e_{\text{stop}}], V = [e_W^{\text{all}} || e_{\text{stop}}], Q = z_W, \text{mask} = m_W) \\
p_A &\leftarrow \text{Attention}(K = e_A^{\text{all}}, V = e_A^{\text{all}}, Q = z_A, \text{mask} = m_A)
\end{aligned}$$

where e_{SG} , e_W , and e_A are predicted embeddings of the crystal’s space group, occupied Wyckoff positions, and atomic elements, respectively; l_{max} is the number of encoder layers; m_{causal} is a causal attention mask enforcing atom orderings; e_W^{all} , e_{stop} , and e_A^{all} are predicted embeddings of all sampleable Wyckoff positions, the stop token, and all sampleable elements, respectively; m_W and m_A are attention masks enforcing lexicographic atom orderings; $p_{W,\text{stop}}$ and p_A are the model probabilities over Wyckoff positions, the stop token, and elements; and $\text{EncoderLayer}(x, m)$ is a module summarized by the following:

$$\begin{aligned}
x_{\text{norm}} &\leftarrow \text{LN}(x) \\
x &\leftarrow x + \text{MultiheadSelfAttention}(x_{\text{norm}}, m) \\
x &\leftarrow \text{Dropout}(x) \\
x &\leftarrow \text{Dropout}(x + \text{Linear}(\text{Dropout}(\text{MLP}(x))))
\end{aligned}$$

A.4.2 Graph Neural Network

Our graph neural network architecture can be summarized as follows. Our GNN was a modified version of FAENet [20], replacing sum pooling with variance-preserving aggregation [81] and removing frame averaging since we trivially achieve $\text{SE}(3)$ invariance by canonicalizing crystals with the ASU representation. We constructed fully connected atom graphs \mathcal{G} wherein each atom in the conventional unit cell was connected to every other atom in the conventional unit cell by their minimum-length edge under periodic boundary conditions. If ties existed, all corresponding edges were included. At initialization, to build features which were invariant to lattice translations, we sampled frequencies $\nu_k \in \mathbb{Z}^3 \cap [-512, 512]^3 \setminus \mathbf{0}$ without replacement with probabilities drawn from

a discretized standard normal distribution. The architecture is summarized as follows:

$$\begin{aligned}
t_{emb} &= \text{MLP}(t) \\
\text{PlaneWaveEmbedding}(\cdot) &\leftarrow [\cos(2\pi(\cdot)\nu_1) \parallel \sin(2\pi(\cdot)\nu_1) \parallel \cos(2\pi(\cdot)\nu_2) \parallel \sin(2\pi(\cdot)\nu_2) \parallel \dots] \\
h_i^0 &\leftarrow \text{MLP}(e_{a_i} \parallel \text{PlaneWaveEmbedding}(x_i) \parallel t_{emb}) \\
e_{ij} &\leftarrow \text{MLP}(\text{PlaneWaveEmbedding}(x_j - x_i) \parallel \text{RBF}(d_{ij}) \parallel \text{Norm}(\mathbf{l})) \\
f_{ij}^l &\leftarrow \text{MLP}(e_{ij} \parallel h_i^l \parallel h_j^l) \\
h_i^{l+1} &\leftarrow h_i^l + a \cdot \text{MLP} \circ \text{GraphNorm} \left(\frac{1}{\sqrt{|\mathcal{N}_i|}} \sum_{j \in \mathcal{N}_i} h_j^l \odot f_{ij}^l \right) \\
h_i^{l_{\max}} &\leftarrow \text{MLP}(h_i^0 \parallel \dots \parallel h_i^{l_{\max}}) \\
\hat{f}_i &\leftarrow \text{MLP}(h_i^{l_{\max}})
\end{aligned}$$

where t is the diffusion timestep; $x_i \in [0, 1]^3$ is the fractional coordinate of atom i ; d_{ij} is the minimum-image pairwise Cartesian distance between atoms i and j , i.e., for lattice matrix $L \in \mathbb{R}^{3 \times 3}$,

$$d_{ij} = \min_{n_1, n_2, n_3 \in \mathbb{Z}} \|Lx_i - L(x_j + n_1 + n_2 + n_3)\|;$$

h_i^l is the node feature of the i th atom after l rounds of message passing; $\mathbf{l} \in \mathbb{R}^6$ are the conventional lattice parameters; l_{\max} is the number of message passing layers; a is a learnable scalar initialized to zero; e_{a_i} is the embedding of atom i 's element; and $\hat{f}_i \in \mathbb{R}^3$ is the non-symmetrized score prediction (see Eq. 3) for atom i . For the LDiff variant of our model, we additionally had

$$\epsilon_L \leftarrow \text{MLP} \left(\frac{1}{N} \sum_{i=1}^N h_i^{l_{\max}} \right)$$

where $\epsilon_L \in \mathbb{R}^6$ is the predicted noise for the lattice k -vector (see [41]).

We see that \hat{f}_i is invariant to lattice translations since the model's spatial inputs are the minimum-image distances d_{ij} under periodic lattice translations and translation invariant plane wave embeddings of relative and absolute atomic fractional coordinates ($x_j - x_i$ and x_i , respectively). We symmetrized \hat{f}_i with space group equivariance according to Equation 3.

A.4.3 Space group and Wyckoff position features

To represent a space group, we created one-hot features indicating the 6 lattice centering types, 6 crystal families, 32 crystallographic point groups, whether the space group is chiral, and whether the space group is centrosymmetric. Since this representation ignores fractional translations from glide and screw symmetries, we counted the maximum number of collisions between space group representations (16), arbitrarily indexed colliding space groups, and created additional one-hot features to avoid collisions. This yielded $6 + 6 + 32 + 1 + 1 + 16 = 62$ features. While the last 16 features are not physically motivated, they give the model the capacity to differentiate space groups and learn their relationships. These features are more abstracted than those of SymmCD [55], which used binary features indicating the 26 space group symmetry operations, 15 symmetry axes, and 7 lattice systems, totaling $7 + 26 \times 15 = 397$ features.

To featurize a Wyckoff position, we used our 62 space group features; one-hot features for the 100 unique site symmetry symbols across space groups, the 4 numbers of degrees of freedom in atom coordinates (0, 1, 2, or 3), and the 17 unique Wyckoff multiplicities; and 48 Fourier features of Wyckoff shapes' midpoints and vertices (see Sec 4.1). This yielded $62 + 100 + 4 + 17 + 48 = 231$ Wyckoff position features. To compute Fourier features $\nu \in \mathbb{R}^{48}$, we collected the midpoint and vertices of each Wyckoff shape belonging to the Wyckoff position. We used these points

$\{p_i \in \mathbb{R}^{3 \times 1}\}_{i=1}^{n_w}$ to compute the following:

$$\begin{aligned} f &= 2\pi \cdot \text{linspace}(\text{start} = 1, \text{end} = 32, \text{steps} = 8) \in \mathbb{R}^{1 \times 8} \\ \nu_i &= \text{flatten}(p_i f) \in \mathbb{R}^{24} \\ \nu &= \frac{1}{n_w} \sum_i^{n_w} [\sin(\nu_i), \cos(\nu_i)] \in \mathbb{R}^{48} \end{aligned}$$

We confirmed that this representation yielded no collisions between featurizations of different Wyckoff positions. Our representation differs from that of SymmCD, which represents each Wyckoff position by one-hot features of the 15 symmetry axes and 13 symmetry operations (totaling $15 \times 13 = 195$ features) corresponding to the Wyckoff position's stabilizer group.

A.5 Ablation results

Table 5: SGEquiDiff Ablations on MP-20. We replaced our space group and Wyckoff position features with those of SymmCD [55] (+*SymmCD features*) and our denoising graph neural network with that of DiffCSP [40], DiffCSP++ [41], and FlowMM [62] (+*CSPNet*). For fair comparison in the latter experiment, we fixed the number of message passing steps (5), the number of Gaussian edge frequencies (96), and the number of trainable parameters (2.2M) by reducing CSPNet's hidden dimension from 512 to 224. Standard deviations over three training runs are reported in parentheses as errors of rightmost digits.

	Validation Wyckoff log-likelihood \uparrow	Validity (%) \uparrow	U.N. rate (%) \uparrow	Distribution distance \downarrow	CMD \downarrow	Diversity \uparrow			
	Structure	Composition	Template	W_p	$W_{N_{eff}}$	$JSD_{d_{\text{Wyckoff}}}$	Structure	Structure	Composition
SGEquiDiff	-0.3206 ₍₈₃₎	99.25 ₍₂₈₎	86.16 ₍₃₂₎	18.6 ₍₁₆₎	80.0 ₍₁₆₎	0.1938 ₍₁₀₄₉₎	0.2092 ₍₆₇₈₎	0.0252 ₍₉₀₎	0.1718 ₍₅₈₎
+ <i>SymmCD features</i>	-0.4925 ₍₃₀₎	98.55 ₍₂₁₎	82.82 ₍₂₈₎	30.0 ₍₂₁₎	84.9 ₍₁₈₎	0.4217 ₍₁₅₂₅₎	0.1706 ₍₂₇₉₎	0.0395 ₍₁₃₈₎	0.3307 ₍₁₂₅₎
+ <i>CSPNet</i>	-	99.00 ₍₅₎	-	-	80.6 ₍₁₈₎	-	-	0.1823 ₍₂₁₎	0.9252 ₍₁₅₅₎

Table 6: LDiff ablation on MP-20. We replaced our autoregressive lattice sampling method with joint diffusion of atom coordinates and E(3)-invariant lattice matrix representations proposed in DiffCSP++ [41] (+*LDiff*). We report metrics of validity for 10,000 sampled crystals (including the percentage of generated lattices with positive determinant, a requirement to enforce symmetry of chiral space groups); uniqueness and novelty for 1,000 valid crystals; diversity and goodness-of-fit distributional distances for valid, unique, and novel crystals; and stability for valid, unique, and novel crystals. Stability is defined as having a DFT-predicted $E_{\text{hull}} < 0.1$ eV/atom with respect to the Materials Project v2022.8.23. Best result is bolded, second best is underlined. Standard deviations over three training runs are reported in parentheses as errors of rightmost digits.

	Time \downarrow (s / batch)	$ L > 0$	Validity (%) \uparrow Structure	U.N. rate (%) \uparrow Composition	Distribution distance \downarrow W_p	CMD \downarrow Structure	Diversity \uparrow Structure	S.U.N. (%) \uparrow Composition	S.U.N. (%) \uparrow MP-2024
SGEquiDiff	226 ₍₁₂₎	100 ₍₀₎	99.25 ₍₂₈₎	86.16 ₍₃₂₎	18.6 ₍₁₆₎	80.0 ₍₁₆₎	0.1938 ₍₁₀₄₉₎	0.2092 ₍₆₇₈₎	0.0252 ₍₉₀₎
+ <i>LDiff</i>	227 ₍₁₂₎	98.95 ₍₄₉₎	97.02 ₍₉₇₎	84.10 ₍₆₃₎	20.3 ₍₂₀₎	80.3 ₍₁₀₎	0.1435 ₍₆₁₃₎	0.1292 ₍₈₆₎	0.1733 ₍₅₃₎

^{*} For efficiency, stability calculations were conducted on 100 random valid, unique, and novel structures.

A.6 Examples of S.U.N. crystals

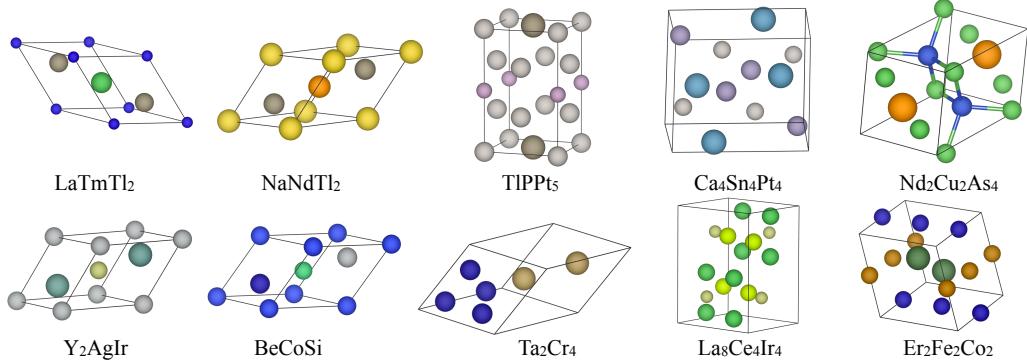


Figure 4: S.U.N. crystals generated by SGEquiDiff trained on MP20.

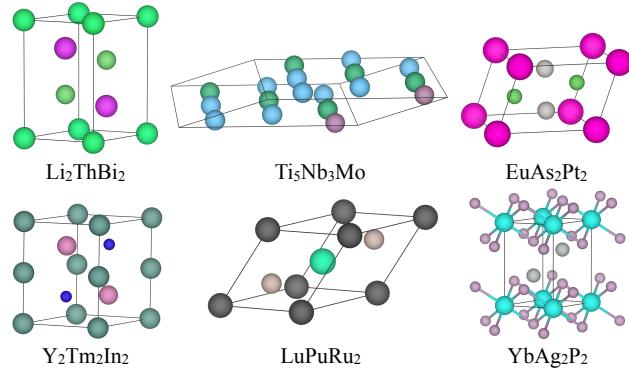


Figure 5: S.U.N. crystals generated by the +LDiff variant of SGEquiDiff trained on MP20.

A.7 Training and hyperparameters

Our code was written with PyTorch [72] and PyTorch Geometric [21]. The model was trained with the Adam optimizer [50] on a single NVIDIA A40 GPU. Training SGEquiDiff took approximately 10 hours on MP20 and 12 hours on MPTS52. Data splits were the same as provided by Xie et al. [100] and Baird et al. [4]. Hyperparameters were tuned manually.

We adopted a modular architecture, simultaneously but independently training the space group sampler, lattice sampler, Wyckoff-Element transformer, and denoising graph neural network. Early stopping based on validation performance was applied to each module separately, eliminating the need for gradient balancing between modules.

Hyperparameter	Value
Batch size	256
Number of epochs	1000
Teacher forced lattice length noise range	0.3
Teacher forced lattice angle noise range	5.0
Lattice sampler hidden dimension	256
Lattice length Fourier scale	5.0
Lattice angle Fourier scale	1.0
Lattice length bin edges Fourier scale	2.0
Lattice angle bin edges Fourier scale	1.0
Transformer dropout rate	0.1
Transformer hidden dimension	256
Transformer number of hidden layers	4
Transformer number of heads	2
SGWN lattice translations	3
SGWN Monte Carlo samples	2500
Time embedding dimension	128
Number of plane wave frequencies	96
Number of Cartesian distance radial basis functions	96
Radial basis functions cutoff distance	10.0
Edge embedder hidden dimension	128
Node embedder hidden dimension	256
Number of message passing steps	5
Space group learning rate	1×10^{-3}
Lattice learning rate	1×10^{-4}
Transformer learning rate	1×10^{-4}
GNN learning rate	1×10^{-3}
Weight decay	0.0
Learning rate scheduler	ReduceLROnPlateau
Scheduler factor	0.6
Scheduler patience	30
Minimum learning rate	1×10^{-5}
Gradient clipping by value	0.5
Number of parameters	5,513,396

A.8 Density Functional Theory Calculations

Total energy calculations were performed with density functional theory (DFT) using the Vienna Ab Initio Simulation Package (VASP) [54] with Projector Augmented Wave (PAW) pseudopotentials [8]. The Perdew-Burke-Ernzerhof (PBE) [73] exchange-correlation functional was used for structural relaxations. All input parameters, plane wave cutoffs, convergence criteria, and k-point densities were determined by `pymatgen MPRelaxSet` [69].

Relaxations which don't converge with default settings from VASP (version 6.4) and `MPRelaxSet` do not necessarily indicate issues with generated crystal structures. Specifically, we encountered and attempted to resolve three failure modes during relaxations: errors related to symmetry-finding, issues converging energy and electronic charge density within DFT's self-consistent field (SCF) loops, and issues converging atomic positions. To resolve symmetry-finding errors (reported as `SGRCON` or `INVGRP` in VASP's output files), we tightened the tolerance of the symmetry detection algorithm by setting `SYMPREC=1E-6` from the default `1E-5`. To resolve SCF convergence issues, we changed the electronic minimization algorithm to the slower but more stable `ALGO=NORMAL` from the default `ALGO=FAST`. If SCF-related issues persisted, we additionally set `AMIX=0.2` from the default `0.4`, which results in slower but more stable convergence of the charge density, and increased the number of electronic optimization steps with `NELM=200` (default 100). Finally, to resolve issues related to convergence of atomic positions, we changed the structure optimization algorithm to the slower but more precise `IBRION=1` (default 2). We emphasize that these changes in VASP parameters will change the relaxation trajectory but not change the minimal energy structure or loosen convergence criteria relative to `MPRelaxSet`'s default settings.

For direct comparison with Materials Project (MP) data, we applied the MP's energy correction scheme for anions and mixing GGA/GGA+ U calculations [37, 38]. Crystal stabilities were determined through convex hull analyses using all competing phases available in MP v2022.8.23 [39].

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: Yes, the claims in the abstract and introduction are reflective of the results in the rest of the paper.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We included a "Limitations" section at the end of the paper.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: We included proofs in our appendix.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: In the appendix, we included details of our model architecture, hyperparameters, and other important information related to our implementation.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: In our abstract, we have provided a link to code, pre-processed data, and model checkpoints. However, software for reproducing density functional theory calculations will not be provided as it requires a paid license to the Vienna Ab initio Simulation Package.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We provided all hyperparameters and training details in the appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We could not put error bars on all reported metrics due to computational constraints, but we do report error bars on RMSDs between as-generated and DFT-relaxed crystal structures in the appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer “Yes” if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We provided information on the type of compute, batch sizes, training times, and inference speeds in the main text and appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We have read and followed the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: We have included a broader impacts statement at the end of our paper.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.

- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: Since the synthesis, scaled up processing, and integration of new crystalline materials into devices is a substantial bottleneck in the materials engineering pipeline, we do not believe there is currently high risk for misuse of crystal generative models.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We have cited authors from which all (open source) baseline models and datasets were sourced.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: Links to code and model checkpoints are provided and documented.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: We did not do any crowdsourcing or research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: We did not do any crowdsourcing or research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.

- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: LLMs were not involved in any important, original, or non-standard components of this research.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.