Space Group Conditional Flow Matching

Omri Puny
Weizmann Institute of Science

Yaron Lipman
FAIR at Meta

Benjamin Kurt Miller
FAIR at Meta

Abstract

Inorganic crystals are periodic, highly symmetric arrangements of atoms in three-dimensional space. Their structures are constrained by the symmetry operations of a crystallographic *space group* and restricted to lie in specific affine subspaces known as *Wyckoff positions*. The frequency an atom appears in the crystal and its rough positioning are determined by its Wyckoff positions. Most generative models that predict atomic coordinates overlook these symmetry constraints, leading to unrealistically high populations of proposed crystals exhibiting limited symmetry. We introduce Space Group Conditional Flow Matching (SGFM), a novel generative framework that samples significantly closer to the target population of highly symmetric, stable crystals. We achieve this by conditioning the entire generation process on a given space group and set of Wyckoff positions; specifically, we define a conditionally symmetric noise distribution and a group-conditioned, equivariant, parametric vector field that restricts the motion of atoms to their initial Wyckoff position. Our form of group-conditioned equivariance is achieved using an efficient reformulation of *group averaging* tailored for symmetric crystals. It reduces the computational overhead of symmetrization to a negligible level. SGFM achieves state-of-the-art performance in *De Novo Generation* on the MP-20 benchmark.

1 Introduction

The properties of many technologies are determined by the quality of the solid-state crystals from which they are derived. Better materials lead to better technology. At the same time, the space of possible crystal structures is vast and underexplored. Discovery of novel crystalline structures is critical for material design and technological progress. We explore generative modeling to propose novel crystals. Crystals are solid materials characterized by a periodic arrangement of their constituent atoms. The crystalline structure can be represented by three components: lattice parameters (the geometry of the repeating unit cell), fractional coordinates (the position of each atom within the cell), and the atomic number of each atom. A crystal's space group, a discrete subgroup of the Euclidean group, describes the symmetry of the crystal. The space group restricts atomic locations and lattice structure. Atomic constraints are called Wyckoff positions, sets of symmetrically equivalent points within a unit cell. The Wyckoff positions of a space group partition the unit cell according to the structure of the orbits induced by the group (see Figure 2 for a 2D example).

Most existing generative methods overlook key crystallographic properties: space group and Wyckoff positions. This limits their ability to generate non-trivially symmetric crystals. The empirical distribution from an experimental and computational database of stable crystals (MP-20 (Jain et al., 2013)) contains only 1.7% of so-called *P1* crystals, i.e., the least symmetric crystals with only translational symmetry. This distribution is not merely a bias in the database; symmetry plays a key role in determining optical, electrical, magnetic, and structural properties of materials Chen et al. (2022); Tang et al. (2019); Malgrange et al. (2014); Yang (2005).

We develop a generative model that samples crystals conditioned on a given space group and associated Wyckoff positions. This approach offers two key benefits: (1) it provides greater control over the structure and symmetry of the generated crystals, and (2) it can leverage the lower-dimensional

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

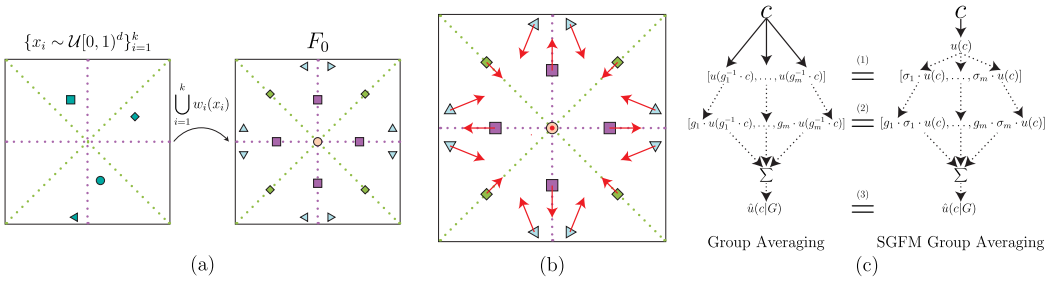


Figure 1: The main components of SGFM. (a) Wyckoff position noise prior. General points are sampled randomly and projected according to the conditioned Wyckoff positions. (b) Space group equivariant vector field. The equivariance of the model combined with the G-symmetry of the input crystal ensures that atoms preserve their symmetry structure. (c) Comparison between standard group averaging and our optimized formulation. The G-symmetry of the input crystal allows space group operations to be replaced by their corresponding permutations (denoted using σ). When combined with the S_n-equivariance of u, this reduces the number of required forward passes per crystal. For visualization purposes, heavy arrows indicate expensive model forward passes.

constraints imposed by Wyckoff positions for improving generation. In contrast to prior methods that incorporate space group information but rely on projection steps to correct atomic placements Jiao et al. (2024); Levy et al. (2025), our model is designed to inherently preserve the assignment of atoms to their designated Wyckoff positions throughout the generation process. This more natural inductive bias translates directly into superior empirical performance.

Our proposed model, *Space Group Conditional Flow Matching* (SGFM), is based on the *Flow Matching* (FM) generative framework Albergo & Vanden-Eijnden (2022); Liu et al. (2022); Lipman et al. (2022). We chose the FM framework for two key reasons: it allows us to use customized source distributions and provides known conditions for generating data that respects specified symmetries Köhler et al. (2020). Based on these advantages, we designed SGFM with two main components: A space group and Wyckoff position conditioned noise prior, which has positive support only for crystal structures that adhere to the symmetry constraints described by the Wyckoff positions; and a group conditioned equivariant vector field, which is a neural network architecture that supports arbitrary space group equivariance. Equivariance is achieved through Group Averaging (GA) Yarotsky (2022), a symmetrization technique that projects arbitrary functions onto their equivariant versions. Although GA is typically computationally expensive, we introduce an efficient formulation tailored for symmetric crystals, reducing the computational overhead of the symmetrization to a negligible level. Our contributions are:

- We define symmetric crystal generation in terms of distributional symmetry properties (Section 3.1), and introduce the conditions that enable flow-based models to sample from such distributions (Section 3.2). We instantiate this flow model as SGFM (Section 3.3), which consists of a noise prior conditional on Wyckoff position along with a space group-equivariant vector field, ensuring that the generated crystals preserve the specified symmetries.
- We propose a novel implementation of GA for symmetric crystals, equivalent in output to the standard GA but *significantly* more efficient (see Figure 1 (c)).
- SGFM achieves state-of-the-art S.U.N. rate on the MP-20 dataset in *de novo generation*.

2 Preliminaries

Equivariance & Invariance. A function $\varphi : \mathcal{X} \to \mathcal{Y}$ is equivariant with respect to a group G if the action of any group element on the input corresponds to a consistent transformation of the output. Equivariance implies $\varphi(g \cdot x) = g \cdot \varphi(x)$ for all $x \in \mathcal{X}$ and $g \in G$. Invariance is a simplified case of equivariance, with all $g \in G$ mapping to a trivial group action on the output space $\varphi(g \cdot x) = \varphi(x)$. Invariance and equivariance also extend to group products: Let $(g_1, g_2) \in G_1 \times G_2$, φ is $G_1 \times G_2$ equivariant if $\varphi((g_1, g_2) \cdot x) = (g_1, g_2) \cdot \varphi(x)$. Additionally, *$G$-invariant distributions* refer to distributions which have a G-invariant density function.

Crystal Representation. A crystal can be represented using the tuple $c' = (\boldsymbol{L}, \boldsymbol{F}, \boldsymbol{A}) \in \mathcal{C}'$, where $\boldsymbol{L} \in \mathbb{R}^{3\times3}$ is a positive-determinant, invertible matrix defining the lengths, angles, and orientation of

the positive-volume unit cell; $\boldsymbol{F} \in [0,1)^{n\times 3}$ denotes the fractional coordinates of n atoms within the unit cell; and $\boldsymbol{A} \in \{0,1\}^{n\times h}$ is a one-hot matrix indicating the atom types in $\boldsymbol{F}$ from a set of h atom types. We adopt the space group–conditioned lattice parameterization proposed by Jiao et al. (2024), which replaces $\boldsymbol{L}$ with a rotation-invariant vector $k \in \mathbb{R}^6$ of coefficients of a symmetric matrix basis, and represent a crystal as $c = (k, \boldsymbol{F}, \boldsymbol{A}) \in \mathcal{C}$. Further details are provided in Appendix C.

Symmetries & Crystals. Our method focuses on how symmetry groups act on crystals. The permutation group S_n acts on c by permuting the rows of $\boldsymbol{F}$ and $\boldsymbol{A}$. If $\sigma \in S_n$ is represented by a permutation matrix $\boldsymbol{P} \in \{0,1\}^{n\times n}$ then $\sigma \cdot c = (k, \boldsymbol{PF}, \boldsymbol{PA})$. The group of isometries of $\mathbb{R}^3$, i.e, the Euclidean group $E(3)$, acts on c by applying an orthogonal transformation $\boldsymbol{R} \in O(3)$ and a translation $\tau \in \mathbb{R}^3$ to the coordinates. For group element $g = (\boldsymbol{R}, \tau) \in E(3)$, the action is defined as $g \cdot c = (k, \boldsymbol{F}\boldsymbol{R}^T + \mathbf{1}_n\tau^T - \lfloor \boldsymbol{F}\boldsymbol{R}^T + \mathbf{1}_n\tau^T \rfloor, \boldsymbol{A})$ where $\mathbf{1}_n \in \{1\}^n$ is a column vector of ones and $\lfloor\cdot\rfloor$ is the element-wise floor function. We specify the action of the product group on c. Let $(g,\sigma) \in G \times S_n$ and $(g,\sigma)\cdot c \coloneqq g \cdot (\sigma \cdot c)$. If $G \leqslant E(3)$ then $g \cdot (\sigma \cdot c) = \sigma \cdot (g \cdot c)$; namely, the operators commute (Puny et al., 2021).

Space Groups. The space group concept formalizes the intrinsic symmetry of a crystal. If $c \in \mathcal{C}$ is a crystal and $G \leqslant E(3)$ is its symmetry space group, then for any $g \in G$ there exists a permutation $\sigma \in S_n$ that satisfies the relation $g \cdot c = \sigma \cdot c$, a property we will denote as *G-symmetry*. In essence, any action of the space group is equivalent to a permutation of the atom positions, Figure 2 visualizes this property using a 2D example. The p4mm symmetry group includes rotations by angles of $\frac{\pi z}{2}$ for $z \in \mathbb{Z}$. The corresponding permutation is invisible, without fabricated labels, because it rearranges the positions of identical shapes. Two Crystals $c_1, c_2 \in \mathcal{C}$ are *Mutually G-Symmetric* if every space group element g corresponds to the same permutation σ on both crystals. Formally, c_1 and c_2 are mutually G-symmetric if $g \cdot c_1 = \sigma \cdot c_1 \iff g \cdot c_2 = \sigma \cdot c_2$. There exist 230 distinct space groups in three-dimensional crystallography. Owing to the intrinsic periodicity of crystal structures, all corresponding subgroups are finite subgroups of the Euclidean group $E(3)$. For non-orthogonal lattice structures, the space group acts on fractional coordinates as elements of the special affine group $SA(3)$, rather than $E(3)$, Figure 4 demonstrates this in 2D. In simpler terms, we apply the group action after mapping every lattice to the cube using $\boldsymbol{L}^{-1}$.

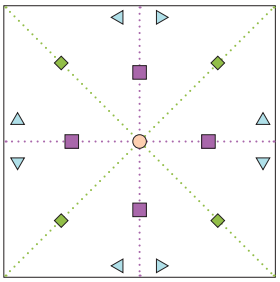


Figure 2: A 2D example of a unit cell with p4mm symmetry. Applying a group element to this set permutes "atoms" of the same type (shape and color). The symmetry divides the unit cell (black box) into four Wyckoff positions: the center, horizontal and vertical coordinate axes, diagonal axes, and general position (white space).

Wyckoff Positions. Intuitively, Wyckoff positions of a space group G indicate regions with specific symmetry properties. Atoms in general position occupy the least symmetric position in the crystal, appear most frequently in the unit cell, and enjoy the fewest restrictions on their coordinates. Meanwhile, atoms in one of the several special positions occupy a region of higher symmetry, appear less frequently in the unit cell, and are restricted to lie in low-dimensional affine subspaces. More formally, a Wyckoff position w is defined by a set of $|w| = m$ affine projections $\{(\boldsymbol{V}_i, \tau_i)\}_{i=1}^m$ onto the corresponding affine subspace, with $\boldsymbol{V}_i \in \mathbb{R}^{3\times 3}$ and $\tau_i \in \mathbb{R}^3$. We denote projection of $f \in [0,1)^3$ onto each of these m affine subspaces by,

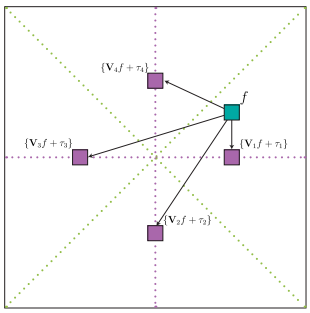


Figure 3: Projection of a coordinate f using Wyckoff position w.

$$w(f) \coloneqq \{\boldsymbol{V}_i f + \tau_i - \lfloor \boldsymbol{V}_i f + \tau_i \rfloor\}_{i=1}^m . \tag{1}$$

If w is a Wyckoff position of a space group G, for every $y \in w(f)$ we have the property $w(f) = G \cdot y$ where $G \cdot y \coloneqq \{g \cdot y \mid g \in G\}$ denotes the *orbit* of y under G. This is visualized in Figure 3; it illustrates how f is mapped to an orbit induced by G through w. Each affine transformation $(\boldsymbol{V}_i, \tau_i)$ identifies with a subgroup $G' \leqslant G$ that fixes points on its corresponding image. Namely, $y_i = \boldsymbol{V}_i f + \tau_i - \lfloor \boldsymbol{V}_i f + \tau_i \rfloor$ for some $f \in [0,1)^3$ if and only if $G' = \{g \in G \mid g \cdot y_i = y_i\}$. That

means that $G' = G_{y_i}$, the *site-symmetry* (stabilizer) group of y_i. We define the crystal c (with fractional coordinates $\boldsymbol{F}$) to be $\mathcal{W}$-*constructable* with respect to $\mathcal{W} = (w_1, \ldots, w_k)$ if there exist k points $\left\{x_i \in [0,1)^3\right\}_{i=1}^{k}$ in the unit cell such that $\boldsymbol{F} = \bigcup_{i=1}^{k} w_i(x_i)$ up to a permutation. In other words, the coordinates in $\boldsymbol{F}$ must be partitionable into complete symmetry orbits, one for each entry in $\mathcal{W}$, with no additional points. This requires the symmetry relations within each orbit, rather than merely requiring individual coordinates to lie on the corresponding Wyckoff subspaces.

Flow Matching (FM) is a generative modeling framework that transforms samples from a simple base distribution p_0 into a complex target distribution p_1 using a time-dependent diffeomorphic map, called a flow, $\psi : [0,1] \times \mathcal{X} \to \mathcal{X}$. This flow is defined through the differential equation:

$$\frac{\mathrm{d}}{\mathrm{d}t}\psi_t(x) = v_t(\psi_t(x)), \quad \psi_0(x) = x \tag{2}$$

where $v_t : [0,1] \times \mathcal{X} \to \mathcal{X}$ is a vector field that governs the evolution of the ψ_t. The flow induces a time-dependent probability density path $p_t : [0,1] \times \mathcal{X} \to \mathbb{R}$ starting at p_0 and ending with p_1. FM trains a parametric approximation u_t of the true vector field v_t Lipman et al. (2024) by solving a regression objective. A conditional flow $\psi(\cdot \,|\, y) : [0,1] \times \mathcal{X} \to \mathcal{X}$ transports the entire base distribution to a single target point $y \in \mathcal{X}$ and is governed by vector field $v_t(\cdot \,|\, y)$, which is easy to compute, unlike $v_t(\cdot)$. Lipman et al. (2022) demonstrated that optimizing u_t to match $v_t(\cdot \,|\, y)$ with regression leads to the same optimum as matching the marginal vector field v_t.

3 Method

This section presents our proposed model, SGFM, a flow matching-based generative approach designed to sample crystal structures conditioned on specified space groups and Wyckoff positions. We start by formalizing the problem and defining the target distribution we aim to sample from in Section 3.1. Then, Section 3.2 outlines sufficient conditions for a flow model to sample from this distribution. Section 3.3 provides a detailed overview of the model, including its key components, the noise prior and vector field. Finally, Section 3.4 presents the training details of SGFM.

3.1 Problem Definition

Given a finite set of crystals $\{c_1, \ldots, c_m\}$, each associated with a space group and Wyckoff positions $(G_i, \mathcal{W}_i)$ and drawn from an unknown target distribution q, our goal is to design a generative model that samples crystals $c \sim p_1$ such that $p_1 \approx q$. To incorporate the structural information encoded by G and $\mathcal{W}$, we factorize the distribution as $p_1(c) = p_1(c \mid G, \mathcal{W})q(G, \mathcal{W})$, enabling us to model the crystal distribution conditionally on G and $\mathcal{W}$. While G and $\mathcal{W}$ are jointly sampled from the empirical distribution, the proposed generative model focuses on sampling crystals from the conditional distribution $c \sim p_1(\cdot \mid G, \mathcal{W})$, which satisfies the following properties:

- $p(\cdot \,|\, G, \mathcal{W})$ is a G-*symmetric* distribution, meaning that $p(c \,|\, G, \mathcal{W}) > 0$ only if c is G-symmetric. The distribution assigns positive probability exclusively to crystals for which G is their corresponding space group.
- $p(\cdot \,|\, G, \mathcal{W})$ is a $\mathcal{W}$-*constructable* distribution, meaning that $p(c \,|\, G, \mathcal{W}) > 0$ only if c is $\mathcal{W}$-constructable. p only supports crystals with fractional coordinates constructible by $\mathcal{W}$.

The definitions provide a full description of the required distribution, but they can be further simplified. The following lemma shows that a $\mathcal{W}$-constructable distribution implies a G-symmetric distribution, where $\mathcal{W}$ denotes the Wyckoff positions associated with the space group G.

Lemma 3.1. *Let $G \leqslant E(3)$ be a space group, $\mathcal{W}$ a corresponding set of Wyckoff positions and $\boldsymbol{F} \in [0,1)^{n\times 3}$ a $\mathcal{W}$-constructable set of points in the unit cell, then $\boldsymbol{F}$ is G-symmetric.*

The proof (Appendix A.3) relies on the fact that the action of a group element on an orbit defines a bijection. Imposing these constraints on the conditional distribution substantially reduces its degrees of freedom, thereby simplifying generation.

3.2 Theoretical Analysis

We now present the theoretical concepts underpinning the development of our generative model. Köhler et al. (2020); Rezende et al. (2019); Song et al. (2023) establish the conditions under which

Algorithm 1 Sample $\boldsymbol{F}_0 \sim p_0(\cdot|G, \mathcal{W})$

Input: $\mathcal{W} = \{w_1, \ldots, w_k\}$ s.t. w_i is a Wyckoff position of the space group G.
Output: $\boldsymbol{F}_0 \in [0,1)^{n\times 3}$ s.t. $n = \sum_{i=1}^{k} |w_i|$.
set $\boldsymbol{F}_0 = [\,]$
for $i = 1$ to k **do**
    Sample $x \sim \mathcal{U}\,[0,1]^3$
    $\boldsymbol{F}_0 = \text{Concatenate}([\boldsymbol{F}_0, w_i(x)])$
end for
return $\boldsymbol{F}_0$

a flow-based model can sample from a G-invariant distribution. For completeness, we restate this theorem and provide a condensed proof in Appendix A.1. Extending these results, we derive the conditions under which a flow-based model can sample from a distribution that is $\mathcal{W}$-constructable. Achieving this requires extending the standard framework with two modifications:

1. Introducing a noise prior that is itself $\mathcal{W}$-constructable.
2. Ensuring that the vector field model u_t is equivariant w.r.t. the space group G and the permutation group S_n. To do so, we extend the vector field (and corresponding flow) equivariance to the group product $G \times S_n$.

Theorem 3.2. *The probability path $p_t(x)$ defined by a flow generated by a $G \times S_n$ equivariant vector field u_t from a $\mathcal{W}$-constructable prior p is $\mathcal{W}$-constructable for all $t \in [0,1]$.*

This theorem follows directly from the lemma below (proof in Appendix A.2), which establishes that if the initial point c_0 of a $G \times S_n$ equivariant flow is a G-symmetric crystal structure, any point along the flow will be mutually G-symmetric with c_0. Based on the mutual G-symmetry, we demonstrate that the flow preserves the site-symmetry structure of c_0. This implies that if c_0 is $\mathcal{W}$-constructable, then $\psi_t(c_0)$ is also $\mathcal{W}$-constructable. Figure 1 (b) visualizes the core idea behind the theorem, illustrating how the equivariant vector field constrains atoms to move solely within the image of their Wyckoff position.

Lemma 3.3. *Let ψ_t be a $G \times S_n$ equivariant flow and $c \in \mathcal{C}$ be G-symmetric and $\mathcal{W}$-constructable, then $\psi_t(c)$ is G-symmetric and $\mathcal{W}$-constructable.*

3.3 SGFM

In this section, we introduce the key components of SGFM, with emphasis on the prior model and the learned vector field architecture. We explain how the previously outlined conditions are concretely implemented. The main focus lies in the interaction between the method and the crystal's fractional coordinates, due to their strong dependence on the space group and Wyckoff positions.

Noise Prior. According to Theorem 3.2, for the generated distribution to satisfy the conditions outlined in Section 3.1, the noise prior must also satisfy the same constraints. Algorithm 1 presents a noise prior sampling pseudocode (with 2D visualizations in Figure 1 (a)) that generates initial fractional coordinates compliant with these requirements. The sampling procedure for the lattice parameters and atom types is described in Section 3.4 for the reader's convenience. The algorithm iterates over the set of Wyckoff positions and samples orbits induced by G by projecting random points from the unit cell. It follows directly from the algorithm's construction that $\boldsymbol{F}_0$ is $\mathcal{W}$-constructable.

The sampling procedure also induces a uniform distribution on each periodic Wyckoff image with respect to its natural measure. Uniform sampling from $[0,1)^3$ corresponds to normalized Haar measure on the torus $\mathbb{T}^3 = \mathbb{R}^3/\mathbb{Z}^3$. Since the linear parts V_i of the Wyckoff maps used in our implementation have integer coefficients, $x \mapsto V_i x \pmod 1$ defines a continuous group homomorphism onto its image and therefore pushes uniform Haar measure to uniform Haar measure on that image. The affine translation τ_i preserves uniformity. This holds for images of any dimension, with zero-dimensional positions yielding point masses.

Space Group Conditional Vector Field. As noted previously, u_t must be $G \times S_n$ equivariant in order for the flow to be G-symmetric and $\mathcal{W}$-constructable. The challenge lies in using a single u_t model across crystals with varying space groups. Since some space groups (with non-orthogonal

lattice structure) act on the fractional coordinates with special affine structured transformations, (see Figure 4), using an $E(3) \times S_n$-equivariant model is non-trivial. In addition, modeling G-equivariance using a supergroup restricts the expressiveness of the model. To address these limitations, we adopt Group Averaging (GA) Yarotsky (2022), a symmetrization operator that projects a backbone model onto the space of G-equivariant functions. It is defined as:

$$\hat{u}(c \,|\, G) = \frac{1}{|G|} \sum_{g \in G} g \cdot u(g^{-1} \cdot c). \tag{3}$$

Applying GA to enforce space group symmetry addresses the previously mentioned limitations. Specifically, if the backbone u is S_n equivariant , then $\hat{u}(\cdot \,|\, G)$ is $G \times S_n$ equivariant Puny et al. (2021). Additionally, GA is not limited to subgroups of $E(3)$ and supports all finite groups. Importantly, Puny et al. (2021) showed that symmetrization preserves the expressive power of the original model. However, a drawback of GA is its computational burden: directly applying Equation (3) increases the number of evaluations of u by a factor of $|G|$, which can be as large as 192 for 3D space groups (the average space group size in the MP-20 dataset is ~ 45). To mitigate this computationally intensive formulation, we leverage the fact that the inputs to $\hat{u}(\cdot \,|\, G)$ are G-symmetric crystals, allowing us to derive an efficient and equivalent formulation of GA specific to this case.

Lemma 3.4. *Let $c \in C$ be a crystal, G its space group, and u an S_n equivariant vector field. Then, Equation* (3) *can be equivalently rewritten as follows:*

$$\hat{u}(c \,|\, G) = \frac{1}{|G|} \sum_{g \in G} g \cdot \sigma_{g^{-1}|c} \cdot u(c) \tag{4}$$

where $\sigma_{g^{-1}|c} \in S_n$ satisfies the equation $\sigma_{g^{-1}|c} \cdot c = g^{-1} \cdot c$.

The formulation presented in Equation (4) requires only a single evaluation of u, which dramatically improves the model efficiency. Figure 1 (c) compares Equation (3) and Equation (4) and visualizes the efficiency gain. Furthermore, computing $\sigma_{g^{-1}|c}$ is computationally efficient, since we can decompose the problem according to the orbits of c, determined by $\mathcal{W}$ (Appendix A.3). At inference time, these permutations only need to be computed once for c_0, since Theorem 3.2 guarantees that the flow preserves G-symmetry structure. During training, permutations are computed only once during preprocessing for every data point. Table 7 compares the training and generation runtime between SGFM, a non-equivariant variant (no symmetrization), and standard GA highlighting the efficiency gains of our GA implementation compared to the standard GA, and demonstrating that its computational cost is comparable to using a backbone without symmetrization. Further details of this comparison are in Appendix I.1.

3.4 Training SGFM

This section provides an overview of the SGFM training process. Let $c_1 \in \mathcal{C}$ be a crystal from the training set with a corresponding space group G and Wyckoff positions $\mathcal{W}$. We will denote $c_0 \sim p_0(\cdot|G, \mathcal{W}, c_1)$ a sample from the conditional noise prior, $c_t = \psi_t(c_0 \,|\, c_1)$ the conditional flow where $c_t = (k_t, \boldsymbol{F}_t, \boldsymbol{A}_t)$, $v_t(c_t|c_1) = (v_t^k(c_t|c_1), v_t^{\boldsymbol{F}}(c_t|c_1), v_t^{\boldsymbol{A}}(c_t|c_1))$ is the conditional vector field and $\hat{u}_t(c_t|G) = (\hat{u}_t^k(c_t|G), \hat{u}_t^{\boldsymbol{F}}(c_t|G), \hat{u}_t^{\boldsymbol{A}}(c_t|G))$ is the prediction of the $G \times S_n$ equivariant vector field parametric model.

Lattice Parameters. As noted in Section 2, we represent lattice parameters using the group-conditioned representation from Jiao et al. (2024), where $k \in \mathbb{R}^6$ encodes the basis coefficients of a 3D symmetric matrix constrained to G-specific subspaces. To sample $k_0 \in \mathbb{R}^6$, we first draw coefficients $k' \sim \mathcal{N}(0, \boldsymbol{I})$ and apply a group condition mask: $k_0 = k' \odot m(G)$, where $m(G) \in \{0, 1\}^6$ is a group-dependent binary mask that zeros out the irrelevant coefficients. For hexagonal lattices we set the first entry of k_0 to be $\frac{-\log(3)}{4}$, as described in Table 3. k_t is computed as a linear interpolation of k_0 and k_1, $k_t = (1-t)k_0 + tk_1$ and the corresponding component of the conditional vector field is $v_t^k(c_t|c_1) = k_1 - k_0$. Since the group action does not directly act on the lattice parameterization we need to apply the mask on $\hat{u}_t^k(c_t|G)$, both in training and after each generation step. The lattice optimization objective, $\mathcal{L}^k(\theta)$, is defined as:

$$\mathbb{E}_{t, q(c_1), p_0(c_0|G)} \left\| \hat{u}_t^k(c_t|G) \odot m(G) - (k_1 - k_0) \right\|_2^2 \tag{5}$$

Table 1: DNG evaluation. Models were trained on the MP-20 dataset. NFE refers to the number of generation steps. We evaluated all the S.U.N. metrics, except those marked with *. *DiffCSP++ (empirical)*, *SGEquiDiff* results were computed with only 1,000 samples; others used 10,000 samples. Methods above the dividing line do not explicitly use Wyckoff positions; methods below use them.

Model	NFE	Validity (%) ↑		Property ↓			Stable (%) ↑ S.U.N. (%) ↑ S.S.U.N. (%) ↑ $E_{\text{hull}} < 100$ meV/Atom			Stable (%) ↑ S.U.N. (%) ↑ S.S.U.N. (%) ↑ $E_{\text{hull}} < 0$ meV/Atom		
		Structural	Composition	d_ρ	d_{elem}	d_{cn}						
CDVAE	5000	**100.00**	86.70	0.688	0.278	-	-	-	-	-	-	-
ADiT	500	99.74	**92.14**	-	-	-	**72.0**	27.4	18.21	13.0	4.6	3.57
FlowMM	500	96.86	83.24	0.075	0.079	0.443	31.2	19.7	14.38	4.6	2.3	1.85
FlowLLM	250	99.81	89.05	0.660	0.090	-	67.9	21.9	20.18	14.2	3.6	3.27
MatterGen	1000	100.00	82.60	0.206	0.242	-	-	24.3	-	-	-	-
OMatG	680	95.05	82.84	**0.060**	**0.017**	0.165	44.4	23.7	13.28	6.6	2.2	1.24
SymmCD	1000	90.34	85.81	0.230	0.400	-	-	-	-	-	-	-
DiffCSP++ (empirical)	1000	99.94	85.12	0.235	0.374	-	31.4	21.1	20.0	7.2	4.0	4.0
DiffCSP++ (WyFormer)	1000	99.66	80.34	0.670	0.098	-	-	-	-	-	3.8*	-
SGEquiDiff	1000	99.25	86.16	0.193	0.209	-	-	24.4*	-	-	-	-
SGFM (empirical)	500	99.87	86.81	0.075	0.181	**0.076**	64.1	**30.3**	**29.03**	**14.6**	**6.9**	**6.7**
SGFM (WyFormer)	500	99.87	84.76	0.237	0.233	-	48.4	22.6	21.72	10.6	4.7	4.53

Atom Types. When we predict atom types, we follow the modeling approach introduced in Miller et al. (2024), where atom types are represented using a $\{-1, 1\}$ binary format instead of standard one-hot encoding. Specifically, $A_1 \in \{0, 1\}^{n \times h}$ is converted into its binary representation $\tilde{A}_1 \in \{-1, 1\}^{n \times \lceil \log_2 h \rceil}$. To ensure G-symmetry in the initial sample c_0, atom types must be consistent within each orbit. Accordingly, we sample Gaussian noise $\mathcal{N}(0, 1)^{\lceil \log_2 h \rceil}$ per orbit and broadcast it to all atoms within that orbit to sample $\tilde{A}_0 \in \mathbb{R}^{n \times \lceil \log_2 h \rceil}$. We define $\tilde{A}_t = (1 - t)\tilde{A}_0 + t\tilde{A}_1$ and $v_t^A(c_t|c_1) = \tilde{A}_1 - \tilde{A}_0$. The atom types optimization objective, $\mathcal{L}^A(\theta)$, is defined as:

$$\mathbb{E}_{t, q(c_1), p_0(c_0|G)} \left\| \hat{u}_t^A(c_t|G) - (\tilde{A}_1 - \tilde{A}_0) \right\|_2^2 \tag{6}$$

Our GA formulation (Equation (4)) ensures that $\hat{u}_t^A(c_t \mid G)$ is G-invariant, meaning the atom type vector field is consistent across orbits, as required. A detailed explanation of this property can be found in Appendix H.3. During inference, we apply the sign function to convert the continuous atom type predictions into their binary representation.

Fractional Coordinates. Algorithm 1 describes a general procedure for sampling fractional coordinates that are both G-symmetric and $\mathcal{W}$-constructable. To ensure G-symmetry of the conditional flow, the initial coordinates $\boldsymbol{F}_0 \sim p_0(\cdot \mid G, \mathcal{W})$ must be G-symmetric with $\boldsymbol{F}_1$. This requires that the order of elements and operators in $\mathcal{W}$ match that used to generate $\boldsymbol{F}_1$, which we precompute during preprocessing using the PyXtal library Fredericks et al. (2021). We adopt the flat torus geometry of the unit cell, following the approach proposed by Miller et al. (2024), and define the conditional flow over the fractional coordinates,

$$\psi_t(\boldsymbol{F}_0|\boldsymbol{F}_1) = \boldsymbol{F}_0 + t \cdot \log_{\boldsymbol{F}_0}(\boldsymbol{F}_1),$$
$$\log_{\boldsymbol{F}_0}(\boldsymbol{F}_1) = \frac{1}{2\pi} \operatorname{atan2}([\sin(\boldsymbol{F}_1 - \boldsymbol{F}_0), \cos(\boldsymbol{F}_1 - \boldsymbol{F}_0)]).$$

Where $\log_{(\cdot)}(\cdot)$ is the element-wise logarithmic map over the flat tori. In Appendix B we demonstrate: (1) the conditional vector field $\log_{\boldsymbol{F}_0}(\boldsymbol{F}_1)$ is G-equivariant but with respect to a different representation of G. Let $g \in G$ then $\log_{g \cdot \boldsymbol{F}_0}(g \cdot \boldsymbol{F}_1) = g^\star \log_{\boldsymbol{F}_0}(\boldsymbol{F}_1)$ where $g^\star$ is defined by the homomorphism $(\boldsymbol{R}, \tau) \mapsto \boldsymbol{R}$; (2) $\psi_t(\boldsymbol{F}_0|\boldsymbol{F}_1)$ is mutually G-symmetric with $\boldsymbol{F}_0$ and $\boldsymbol{F}_1$, hence G-symmetric and $\mathcal{W}$-constructable. The fractional coordinates optimization objective, $\mathcal{L}^F(\theta)$, is defined as:

$$\mathbb{E}_{t, q(c_1), p_0(c_0|G, \mathcal{W})} \left\| \hat{u}_t^{\boldsymbol{F}}(c_t|G) - \log_{\boldsymbol{F}_0}(\boldsymbol{F}_1) \right\|_2^2 \tag{7}$$

Combining all components and hyperparameters $\lambda_k, \lambda_{\boldsymbol{F}}, \lambda_{\boldsymbol{A}} \in \mathbb{R}^+$ we obtain our training objective:

$$\mathcal{L}^{\text{SGFM}}(\theta) = \lambda_k \mathcal{L}^k(\theta) + \lambda_{\boldsymbol{F}} \mathcal{L}^{\boldsymbol{F}}(\theta) + \lambda_{\boldsymbol{A}} \mathcal{L}^{\boldsymbol{A}}(\theta). \tag{8}$$

4 Experiments

Our experiments are divided into two distinct tasks. First, we evaluate SGFM on *De Novo Generation* (DNG). The objective of this task is to simulate a materials discovery campaign by generating atom types, fractional coordinates, and lattice parameters from scratch. Second, to isolate the specific

impact of our model on the relevant crystal structure components, we evaluate it on *Crystal Structure Prediction* (CSP). By providing the model with conditioned atom types and requiring it to generate only the lattice parameters and fractional coordinates, we can more precisely measure how effectively SGFM utilizes Wyckoff positions and space group information, with performance determined by how closely the output matches a specific target crystal with shared atom types, space group and Wyckoff positions. While assuming known atom types and Wyckoff positions is impractical in applications, it serves as a rigorous "unit test" for the generative process.

Metrics. In DNG, performance is measured by the model's ability to produce structures that are simultaneously stable, unique, and novel (S.U.N.). Stable inorganic crystals typically have non-*P1* space groups, according to experimental databases, due to symmetries inherent in the electron density in crystalline matter. We therefore additionally report the symmetric, stable, unique, and novel rate, i.e., S.S.U.N. (Kazeev et al., 2025) which is the S.U.N. rate with an additional restriction to crystals that *do not* have *P1* symmetry. *P1* is the trivial symmetry in periodic systems containing only lattice translations. Given that stability evaluation via DFT (details in Appendix D) is computationally expensive, we include standard proxy metrics, such as *structural & compositional validity* and various property divergences, to further validate the quality of the generated crystals. For CSP, we utilize `StructureMatcher` (Ong et al., 2013) to identify Match Rate (MR) and the root mean square error (RMSE). We use standard CSP settings. Further details are provided in Appendix F.

Datasets. We evaluate DNG on *MP*-20 Jain et al. (2013), with 45,231 diverse crystals from the Materials Project. We perform CSP on an additional four datasets: *MPTS*-52, a time-ordered variant with $40,476$ crystals featuring larger unit cells; *Perov*-5 Castelli et al. (2012), with $18,928$ perovskites sharing a common structure but varying atom types, *Carbon*-24 Pickard (2020), containing $10,153$ carbon crystals with diverse structures, and *Alex-MP*-20, a large-scale set of $607,684$ crystals combining MP-20 and Alexandria data Schmidt et al. (2022a,b).

Baselines. We compared SGFM to several state-of-the-art baselines. Methods that do not incorporate space group information in their generation process include *CDVAE* Xie et al. (2021), *ADiT* Joshi et al. (2025), *FlowMM* Miller et al. (2024), *FlowLLM* Sriram et al. (2024) , *MatterGen* Zeni et al. (2023), DiffCSP Jiao et al. (2023), Sriram et al. (2024), and *OMatG* Hoellmer et al. (2025). In contrast, *SymmCD* Levy et al. (2025), *DiffCSP++* Jiao et al. (2024), *WyFormer* Kazeev et al. (2025), and *SGEquiDiff* Chang et al. (2025) explicitly incorporate space group and Wyckoff position information. Additional details on each baseline are provided in Appendix G.

Model Details. To model $\hat{u}_t$, we adopt the architecture used in Jiao et al. (2023), which utilizes a modified version of *EGNN* (Satorras et al., 2022) to handle fractional coordinates. The model applies sinusoidal embeddings to the fractional coordinates, ensuring invariance to lattice translations in addition to the space group equivariance. A description of the architecture and the hyperparameters used in each experiment are provided in Appendix H. For improved sampling quality, we apply inference anti-annealing Yim et al. (2023) that adjusts the prediction velocity during generation.

De Novo Generation

We evaluated SGFM on MP-20. We employed two configurations of incorporating Wyckoff positions, one using Wyckoff positions taken from the train set, denoted *empirical*, and one from the output of *WyFormer*. All models were evaluated with $10,000$ samples, except *DiffCSP++ (empirical)* and *SGEquiDiff* that used only 1000 samples. *DiffCSP++ (WyFormer)* was reported by the authors using slightly different density functional theory settings. Wasserstein divergence for properties is computed between the test set and a structurally and compositionally valid subset of 1000 generated samples. Training and sampling details are in Appendix H.2.

Results for the MP-20 evaluation are summarized in Table 1. SGFM, utilizing empirical Wyckoff positions, achieves state-of-the-art S.U.N. and S.S.U.N. rates across all energy thresholds. The comparison with SGFM (WyFormer) quantifies the effect of replacing empirical Wyckoff sampling with learned generation. Although empirical sampling yields stronger results, SGFM (WyFormer) still outperforms all compared baselines in S.U.N. and S.S.U.N. at the strict $E_{\text{hull}} < 0$ meV/atom threshold. This demonstrates that SGFM remains effective without access to the empirical Wyckoff distribution, while the performance gap between the two settings suggests room for improvement in learned Wyckoff generation. To further compare SGFM's generation efficiency with other models, we measured generation times (across multiple batch sizes) on a single NVIDIA L40S GPU. The results are reported in Table 6. Despite requiring more parameters (16.2M compared to 12.2M

for DiffCSP++ and 5.5M for SGEquiDiff) and operating on the full conventional cell rather than the asymmetric unit (which reduces memory-footprint), SGFM achieves comparable efficiency to SGEquiDiff and outperforms DiffCSP++.

Table 2: Crystal Structure Prediction. MR denotes match rate. Best results in bold. Models below the dashed line receive ground-truth space-group and Wyckoff information; comparisons across the line therefore involve different conditioning information.

Model	Wyckoff Positions	MP-20		MPTS-52		Perov-5		Carbon-24		Alex-MP-20	
		MR (%) ↑	RMSE ↓	MR (%) ↑	RMSE ↓	MR (%) ↑	RMSE ↓	MR (%) ↑	RMSE ↓	MR (%) ↑	RMSE ↓
CDVAE	None	33.90	.1045	5.34	.2106	45.31	.1138	17.09	.2969	-	-
DiffCSP	None	51.49	.0631	12.19	.1786	52.02	.0760	17.54	.2759	-	-
FlowMM	None	61.39	.0566	17.54	.1726	53.15	.0992	**23.47**	.4122	-	-
OMatG	None	**69.83**	.0741	**27.38**	.1970	**83.06**	.3753	-	-	72.50	.1260
DiffCSP++	Ground truth	80.27	.0295	46.29	.0896	98.44	.0430	-	-	83.18	.0188
SGFM	Ground truth	**82.74**	.0288	**51.79**	.0827	**98.57**	.0188	**55.02**	.0952	**84.40**	.0198

Crystal Structure Prediction

The generative task in CSP requires sampling from the conditional target distribution $c \sim q(\cdot|\boldsymbol{A})$, where $\boldsymbol{A}$ denotes a predefined atom type composition. This conditioning implies that during both training and generation $\boldsymbol{A}_t = \boldsymbol{A}$ for all $t \in [0, 1]$, effectively ignoring the loss term $\mathcal{L}^{\boldsymbol{A}}$ and the atom type component $\hat{u}_t^A(c_t|G)$. For evaluation, a crystal structure is generated for each entry in the test set and then compared against the corresponding ground truth structure.

For the results in Table 2, SGFM and DiffCSP++ receive ground-truth space-group and Wyckoff information. This provides substantial additional information: for Wyckoff positions with zero degrees of freedom, the fractional coordinates are fully determined. We therefore treat this setting as a controlled evaluation of symmetry-conditioned generation, with SGFM and DiffCSP++ forming the appropriate matched comparison. SGFM achieves higher match rates on all four datasets for which DiffCSP++ results are available, while Figure 5 illustrates its sampling-efficiency gains. Results for methods without explicit symmetry conditioning provide useful context, but comparisons must account for the additional information supplied to SGFM and DiffCSP++.

We performed an ablation study to isolate the individual contributions of the SGFM components. Table 9 presents a CSP evaluation on the MP-20 dataset across two key dimensions: the prior distribution (Wyckoff-conditioned vs. uniform over the unit cell) and the vector-field architecture (equivariant vs. non-equivariant). The results indicate that while a Wyckoff-conditioned prior provides a superior initialization, its synergy with the space-group equivariant vector field is the primary driver of improved performance, a finding consistent with our theoretical framework.

To evaluate our method in settings where Wyckoff positions are unknown, we utilize *CSPML* (Kusaba et al., 2022), a metric learning-based model. Given an atomic composition, CSPML retrieves a similar composition from a template set along with associated space group and Wyckoff positions. We use these as inputs for the benchmark reported in Table 8. The results indicate heuristic approaches for sampling Wyckoff positions provide useful information, but less than the ground truth positions.

5 Related Work

There is a growing body of literature about generative models for inorganic crystals. We focus here on works with similar inductive biases, namely explicit utilization of Wyckoff positions. We first consider works that generate atomic coordinates. Cao et al. (2024) created an autoregressive model that generates crystals sequentially in Wyckoff position's lexicographic order. Jiao et al. (2024); Levy et al. (2025) produced diffusion models that both represent crystals within the asymmetric unit, a memory-efficient formulation that contains just one representative per orbit. Neither of these methods utilizes space group equivariance and both require projection steps to keep atomic coordinates within the target Wyckoff positions. A concurrently developed diffusion model by Chang et al. (2025) utilizes the asymmetric unit and, like our method, implements space group equivariance via group averaging without needing projection. However, working in the asymmetric unit does not allow for our efficient reformulation in Equation (4) since representing the crystal structure through the asymmetric unit does not yield a G-symmetric representation. Furthermore, keeping track of which

atoms are in the asymmetric unit introduces some complexity in the implementation. As written, Levy et al. (2025) do not address the crystal structure prediction problem. There is also a class of models that ignore explicit atomic coordinates and generate coarse-grained Wyckoff positions alone, inspired by regression methods Goodall & Lee (2020); Goodall et al. (2022). Zhu et al. (2024); Kazeev et al. (2025); Kelvinius et al. (2025) all take this approach. These models synergize with ours and can generate Wyckoff positions for SGFM to use in DNG in Section 4. Innovations in these methods, if conditional on atom types, could replace CSPML for CSP. Further discussion of other relevant work is left for Appendix E.

6 Conclusions

In this work, we introduced SGFM, an FM-based generative model for crystal structures, conditioned on space group and Wyckoff positions. By design, SGFM produces crystals that satisfy symmetry constraints, relying on sufficient conditions we formulated over the noise prior and vector field. We also implemented an efficient group averaging method, enabling the incorporation of space group equivariance into the vector field model with minimal overhead. Evaluated on the DNG task, SGFM achieved state-of-the-art performance. Future directions include extending the model to an unconditional generation setting, where space group and Wyckoff positions are also generated rather than specified. Another promising direction is to integrate property predictors with SGFM to guide generation toward desired physical properties while preserving crystallographic symmetry.

Acknowledgments and Disclosure of Funding

Research was partially supported by the Israeli Council for Higher Education (CHE) via the Weizmann Data Science Research Center and by research grants from the Estates of Tully and Michele Plesser and the Anita James Rosen and Harry Schutzman Foundations.

References

Albergo, M. S. and Vanden-Eijnden, E. Building normalizing flows with stochastic interpolants. *arXiv preprint arXiv:2209.15571*, 2022.

Albergo, M. S., Boffi, N. M., and Vanden-Eijnden, E. Stochastic interpolants: A unifying framework for flows and diffusions, 2023. URL https://arxiv.org/abs/2303.08797.

Atsushi Togo, K. S. and Tanaka, I. Spglib: a software library for crystal symmetry search. *Sci. Technol. Adv. Mater., Meth.*, 4(1):2384822–2384836, 2024. doi: 10.1080/27660400.2024.2384822. URL https://doi.org/10.1080/27660400.2024.2384822.

Barroso-Luque, L., Shuaibi, M., Fu, X., Wood, B. M., Dzamba, M., Gao, M., Rizvi, A., Zitnick, C. L., and Ulissi, Z. W. Open materials 2024 (omat24) inorganic materials dataset and models. *arXiv preprint arXiv:2410.12771*, 2024.

Cao, Z., Luo, X., Lv, J., and Wang, L. Space group informed transformer for crystalline materials generation, 2024. URL https://arxiv.org/abs/2403.15734.

Castelli, I. E., Landis, D. D., Thygesen, K. S., Dahl, S., Chorkendorff, I., Jaramillo, T. F., and Jacobsen, K. W. New cubic perovskites for one-and two-photon water splitting using the computational materials repository. *Energy & Environmental Science*, 5(10):9034–9043, 2012.

Chang, R., Pak, A., Guerra, A., Zhan, N., Richardson, N., Ertekin, E., and Adams, R. P. Space group equivariant crystal diffusion. *arXiv preprint arXiv:2505.10994*, 2025.

Chen, L., Setty, C., Hu, H., Vergniory, M. G., Grefe, S. E., Fischer, L., Yan, X., Eguchi, G., Prokofiev, A., Paschen, S., Cano, J., and Si, Q. Topological semimetal driven by strong correlations and crystalline symmetry. *Nature Physics*, 18(11):1341–1346, September 2022. ISSN 1745-2481. doi: 10.1038/s41567-022-01743-4. URL http://dx.doi.org/10.1038/s41567-022-01743-4.

Chen, R. T. Q. and Lipman, Y. Flow matching on general geometries, 2024. URL https://arxiv.org/abs/2302.03660.

Davies, D. W., Butler, K. T., Jackson, A. J., Skelton, J. M., Morita, K., and Walsh, A. Smact: Semiconducting materials by analogy and chemical theory. *Journal of Open Source Software*, 4 (38):1361, 2019.

Fredericks, S., Parrish, K., Sayre, D., and Zhu, Q. Pyxtal: A python library for crystal structure generation and symmetry analysis. *Computer Physics Communications*, 261:107810, April 2021. ISSN 0010-4655. doi: https://doi.org/10.1016/j.cpc.2020.107810. URL http://www.sciencedirect.com/science/article/pii/S0010465520304057.

Glass, C. W., Oganov, A. R., and Hansen, N. Uspex—evolutionary crystal structure prediction. *Computer physics communications*, 175(11-12):713–720, 2006.

Goodall, R. E. and Lee, A. A. Predicting materials properties without crystal structure: Deep representation learning from stoichiometry. *Nature Communications*, 11(1):1–9, 2020.

Goodall, R. E., Parackal, A. S., Faber, F. A., Armiento, R., and Lee, A. A. Rapid discovery of stable materials by coordinate-free coarse graining. *Science Advances*, 8(30):eabn4117, 2022.

Hoellmer, P., Egg, T., Martirossyan, M. M., Fuemmeler, E., Shui, Z., Gupta, A., Prakash, P., Roitberg, A., Liu, M., Karypis, G., Transtrum, M., Hennig, R. G., Tadmor, E. B., and Martiniani, S. Open materials generation with stochastic interpolants, 2025. URL https://arxiv.org/abs/2502.02582.

Jain, A., Ong, S. P., Hautier, G., Chen, W., Richards, W. D., Dacek, S., Cholia, S., Gunter, D., Skinner, D., Ceder, G., and Persson, K. A. The Materials Project: A materials genome approach to accelerating materials innovation. *APL Materials*, 1(1):011002, 07 2013. ISSN 2166-532X. doi: 10.1063/1.4812323. URL https://doi.org/10.1063/1.4812323.

Jiao, R., Huang, W., Lin, P., Han, J., Chen, P., Lu, Y., and Liu, Y. Crystal structure prediction by joint equivariant diffusion. *arXiv preprint arXiv:2309.04475*, 2023. URL https://arxiv.org/abs/2309.04475.

Jiao, R., Huang, W., Liu, Y., Zhao, D., and Liu, Y. Space group constrained crystal generation, 2024. URL https://arxiv.org/abs/2402.03992.

Joshi, C. K., Fu, X., Liao, Y.-L., Gharakhanyan, V., Miller, B. K., Sriram, A., and Ulissi, Z. W. All-atom diffusion transformers: Unified generative modelling of molecules and materials, 2025. URL https://arxiv.org/abs/2503.03965.

Kazeev, N., Nong, W., Romanov, I., Zhu, R., Ustyuzhanin, A., Yamazaki, S., and Hippalgaonkar, K. Wyckoff transformer: Generation of symmetric crystals, 2025. URL https://arxiv.org/abs/2503.02407.

Kelvinius, F. E., Andersson, O. B., Parackal, A. S., Qian, D., Armiento, R., and Lindsten, F. Wyckoffdiff–a generative diffusion model for crystal symmetry. *arXiv preprint arXiv:2502.06485*, 2025.

Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Kirklin, S., Saal, J. E., Meredig, B., Thompson, A., Doak, J. W., Aykol, M., Rühl, S., and Wolverton, C. The open quantum materials database (oqmd): assessing the accuracy of dft formation energies. *npj Computational Materials*, 1(1):1–15, 2015.

Köhler, J., Klein, L., and Noé, F. Equivariant flows: Exact likelihood generative learning for symmetric densities. In *Proceedings of the 37th International Conference on Machine Learning*, ICML'20. JMLR.org, 2020.

Köhler, J., Invernizzi, M., De Haan, P., and Noé, F. Rigid body flows for sampling molecular crystal structures. In *International Conference on Machine Learning*, pp. 17301–17326. PMLR, 2023.

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

Kusaba, M., Liu, C., and Yoshida, R. Crystal structure prediction with machine learning-based element substitution. *Computational Materials Science*, 211:111496, 2022.

Levy, D., Panigrahi, S. S., Kaba, S.-O., Zhu, Q., Lee, K. L. K., Galkin, M., Miret, S., and Ravanbakhsh, S. Symmcd: Symmetry-preserving crystal generation with diffusion models, 2025. URL https://arxiv.org/abs/2502.03638.

Lin, P., Chen, P., Jiao, R., Mo, Q., Jianhuan, C., Huang, W., Liu, Y., Huang, D., and Lu, Y. Equivariant diffusion for crystal structure prediction. In *Forty-first International Conference on Machine Learning*, 2024. URL https://openreview.net/forum?id=VRv8KjJNuj.

Lipman, Y., Chen, R. T., Ben-Hamu, H., Nickel, M., and Le, M. Flow matching for generative modeling. In *The Eleventh International Conference on Learning Representations*, 2022.

Lipman, Y., Havasi, M., Holderrieth, P., Shaul, N., Le, M., Karrer, B., Chen, R. T. Q., Lopez-Paz, D., Ben-Hamu, H., and Gat, I. Flow matching guide and code, 2024. URL https://arxiv.org/abs/2412.06264.

Liu, X., Gong, C., and Liu, Q. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022.

Malgrange, C., Ricolleau, C., and Schlenker, M. *Symmetry and physical properties of crystals*. Springer, 2014.

Merchant, A., Batzner, S., Schoenholz, S. S., Aykol, M., Cheon, G., and Cubuk, E. D. Scaling deep learning for materials discovery. *Nature*, pp. 1–6, 2023.

Meredig, B. and Wolverton, C. A hybrid computational–experimental approach for automated crystal structure solution. *Nature materials*, 12(2):123–127, 2013.

Michel, K. J. Symmetry building monte carlo-based crystal structure prediction. *Computer Physics Communications*, 185(5):1389–1393, 2014.

Mila AI4Science, ., Hernandez-Garcia, A., Duval, A., Volokhova, A., Bengio, Y., Sharma, D., Carrier, P. L., Koziarski, M., and Schmidt, V. Crystal-GFN: sampling crystals with desirable properties and constraints. *AI for Accelerated Materials Design Workshop (NeurIPS)*, 2023.

Miller, B. K., Chen, R. T. Q., Sriram, A., and Wood, B. M. Flowmm: Generating materials with riemannian flow matching, 2024. URL https://arxiv.org/abs/2406.04713.

Ong, S. P., Richards, W. D., Jain, A., Hautier, G., Kocher, M., Cholia, S., Gunter, D., Chevrier, V. L., Persson, K. A., and Ceder, G. Python materials genomics (pymatgen): A robust, open-source python library for materials analysis. *Computational Materials Science*, 68:314–319, 2013.

Perdew, J. P., Burke, K., and Ernzerhof, M. Generalized gradient approximation made simple. *Physical review letters*, 77(18):3865, 1996.

Pickard, C. J. Airss data for carbon at 10gpa and the c+n+h+o system at 1gpa. *Materials Cloud Archive*, 2020.0026/v1, 2020. doi: 10.24435/materialscloud:2020.0026/v1.

Pickard, C. J. and Needs, R. Ab initio random structure searching. *Journal of Physics: Condensed Matter*, 23(5):053201, 2011.

Puny, O., Atzmon, M., Ben-Hamu, H., Smith, E. J., Misra, I., Grover, A., and Lipman, Y. Frame averaging for invariant and equivariant network design. *Tenth International Conference on Learning Representations (ICLR)*, 2021.

Rezende, D. J., Racanière, S., Higgins, I., and Toth, P. Equivariant hamiltonian flows, 2019. URL https://arxiv.org/abs/1909.13739.

Riebesell, J., Goodall, R. E., Jain, A., Benner, P., Persson, K. A., and Lee, A. A. Matbench discovery an evaluation framework for machine learning crystal stability prediction. *arXiv preprint arXiv:2308.14920*, 2023.

Saal, J. E., Kirklin, S., Aykol, M., Meredig, B., and Wolverton, C. Materials design and discovery with high-throughput density functional theory: the open quantum materials database (oqmd). *Jom*, 65(11):1501–1509, 2013.

Satorras, V. G., Hoogeboom, E., and Welling, M. E(n) equivariant graph neural networks, 2022. URL https://arxiv.org/abs/2102.09844.

Schmidt, J., Hoffmann, N., Wang, H.-C., Borlido, P., Carriço, P. J., Cerqueira, T. F., Botti, S., and Marques, M. A. Large-scale machine-learning-assisted exploration of the whole materials space. *arXiv preprint arXiv:2210.00579*, 2022a.

Schmidt, J., Wang, H.-C., Cerqueira, T. F., Botti, S., and Marques, M. A. A dataset of 175k stable and metastable materials calculated with the pbesol and scan functionals. *Scientific Data*, 9(1):64, 2022b.

Song, Y., Gong, J., Xu, M., Cao, Z., Lan, Y., Ermon, S., Zhou, H., and Ma, W.-Y. Equivariant flow matching with hybrid probability transport, 2023. URL https://arxiv.org/abs/2312.07168.

Sriram, A., Miller, B. K., Chen, R. T. Q., and Wood, B. M. Flowllm: Flow matching for material generation with large language models as base distributions, 2024. URL https://arxiv.org/abs/2410.23405.

Tang, F., Po, H. C., Vishwanath, A., and Wan, X. Comprehensive search for topological materials using symmetry indicators. *Nature*, 566(7745):486–489, February 2019. ISSN 1476-4687. doi: 10.1038/s41586-019-0937-5. URL http://dx.doi.org/10.1038/s41586-019-0937-5.

Wang, H.-C., Botti, S., and Marques, M. A. Predicting stable crystalline compounds using chemical similarity. *npj Computational Materials*, 7(1):12, 2021.

Wirnsberger, P., Papamakarios, G., Ibarz, B., Racanière, S., Ballard, A. J., Pritzel, A., and Blundell, C. Normalizing flows for atomic solids. *Machine Learning: Science and Technology*, 3(2):025009, 2022.

Wood, B. M., Dzamba, M., Fu, X., Gao, M., Shuaibi, M., Barroso-Luque, L., Abdelmaqsoud, K., Gharakhanyan, V., Kitchin, J. R., Levine, D. S., Michel, K., Sriram, A., Cohen, T., Das, A., Rizvi, A., Sahoo, S. J., Ulissi, Z. W., and Zitnick, C. L. Uma: A family of universal models for atoms, 2025. URL https://arxiv.org/abs/2506.23971.

Wu, Z., Ramsundar, B., Feinberg, E. N., Gomes, J., Geniesse, C., Pappu, A. S., Leswing, K., and Pande, V. Moleculenet: a benchmark for molecular machine learning. *Chemical science*, 9(2): 513–530, 2018.

Xie, T., Fu, X., Ganea, O.-E., Barzilay, R., and Jaakkola, T. S. Crystal diffusion variational autoencoder for periodic material generation. In *International Conference on Learning Representations*, 2021. URL https://arxiv.org/abs/2110.06197.

Yan, K., Saxton, A., Qian, X., Qian, X., and Ji, S. A space group symmetry informed network for o(3) equivariant crystal tensor prediction, 2024. URL https://arxiv.org/abs/2406.12888.

Yan, K., Li, X., Ling, H., Ashen, K., Edwards, C., Arróyave, R., Zitnik, M., Ji, H., Qian, X., Qian, X., and Ji, S. Invariant tokenization of crystalline materials for language model enabled generation, 2025. URL https://arxiv.org/abs/2503.00152.

Yang, J. *An introduction to the theory of piezoelectricity*. Springer, 2005.

Yang, M., Cho, K., Merchant, A., Abbeel, P., Schuurmans, D., Mordatch, I., and Cubuk, E. D. Scalable diffusion for materials generation. *arXiv preprint arXiv:2311.09235*, 2023.

Yarotsky, D. Universal approximations of invariant maps by neural networks. *Constructive Approximation*, 55(1):407–474, 2022.

Yim, J., Campbell, A., Foong, A. Y., Gastegger, M., Jiménez-Luna, J., Lewis, S., Satorras, V. G., Veeling, B. S., Barzilay, R., Jaakkola, T., et al. Fast protein backbone generation with se (3) flow matching. *arXiv preprint arXiv:2310.05297*, 2023.

Zeni, C., Pinsler, R., Zügner, D., Fowler, A., Horton, M., Fu, X., Shysheya, S., Crabbé, J., Sun, L., Smith, J., et al. Mattergen: a generative model for inorganic materials design. *arXiv preprint arXiv:2312.03687*, 2023. URL https://arxiv.org/abs/2312.03687.

Zhu, R., Nong, W., Yamazaki, S., and Hippalgaonkar, K. Wycryst: Wyckoff inorganic crystal generator framework. *Matter*, 7(10):3469–3488, 2024.

A Proofs

A.1 Theorem A.1

Theorem A.1. *The probability path $p_t(x)$ defined by a flow generated by a G-equivariant vector field u_t from a G-invariant prior p is G-invariant for all $t \in [0, 1]$.*

Proof. The proof has two main parts. First, we will show that the flow ψ_t defined by the G-equivariant vector field u_t is G-equivariant. Then, we will use this property to demonstrate that the resulting probability path p_t is G-invariant. As a reminder, the flow $\psi : [0, 1] \times \mathcal{X} \to \mathcal{X}$ is defined by the following ODE:

$$\frac{d}{dt}\psi_t(x) = u_t(\psi_t(x)) \tag{9}$$

$$\psi_0(x) = x \tag{10}$$

To demonstrate that ψ_t is equivariant, we will show that two functions, $\varphi_t(x) \coloneqq \psi_t(g \cdot x)$ and $\phi_t(x) = g \cdot \psi_t(x)$ (for arbitrary $g \in G$) satisfy the same ODE with identical initial conditions.

$$\frac{d}{dt}\varphi_t(x) = \frac{d}{dt}\psi_t(g \cdot x) = u_t(\psi_t(g \cdot x)) = u_t(\varphi_t(x))$$
$$\varphi_0(x) = \psi_0(g \cdot x) = g \cdot x$$

$$\frac{d}{dt}\phi_t(x) = \frac{d}{dt}g \cdot \psi_t(x) = g \cdot \frac{d}{dt}\psi_t(x) = g \cdot u_t(\psi_t(x)) = u_t(g \cdot \psi_t(x)) = u_t(\phi_t(x))$$
$$\phi_0(x) = g \cdot \psi_0(x) = g \cdot x$$

The fourth equality uses the G-equivariance of u_t. We can therefore conclude $\psi_t(g \cdot x) = g \cdot \psi_t(x)$ for every $x \in \mathcal{X}, g \in G$ and $t \in [0, 1]$, which proves that ψ_t is G-equivariant.

It remains to show that p_t defines an invariant probability path.

$$\begin{aligned} p_t(g \cdot x) &= p_0(\psi_t^{-1}(g \cdot x)) \det \left[\frac{\partial \psi_t^{-1}}{\partial x}(g \cdot x) \right] \\ &= p_0(\psi_t^{-1}(x)) \det \left[\frac{\partial \psi_t^{-1}}{\partial x}(g \cdot x) \right] \\ &= p_0(\psi_t^{-1}(x)) \det \left[g \cdot \frac{\partial \psi_t^{-1}}{\partial x}(x) \cdot g^{-1} \right] \\ &= p_0(\psi_t^{-1}(x)) \det \left[\frac{\partial \psi_t^{-1}}{\partial x}(x) \right] = p_t(x) \end{aligned}$$

The second equality follows from the G-equivariance of ψ_t^{-1} and the G-invariance of p_0. The third equality is a consequence of the definition of the Jacobian matrix for equivariant functions, and the final equality relies on standard properties of the determinant.

A.2 Proof of Lemma 3.3

Proof. The first part of the proof, which involves showing that ψ_t is G-symmetric, is straightforward and follows directly from the equivariance properties of ψ_t. Since ψ_t is $G \times S_n$-equivariant (and the group actions commute), it is straightforward to see that it is equivariant to each of the groups separately. Let $g \in G$. Then:

$$g \cdot \psi_t(c) = \psi_t(g \cdot c) = \psi_t(\sigma \cdot c) = \sigma \cdot \psi_t(c)$$

Where the first equality follows from the G-equivariance of ψ_t, the second holds because c is G-symmetric, and the final equality follows from the S_n-equivariance of ψ_t. From the above equation, we also conclude that c and $\psi_t(c)$ are mutually G-symmetric. Now, let $g' \in G_{f_i}$, meaning that g' belongs to the site-symmetry group of f_i, the i^{th} fractional coordinate of c. Since $g' \in G$ there exists

a permutation $\sigma' \in S_n$ s.t $g' \cdot c = \sigma' \cdot c$. Moreover, because $g' \in G_{f_i}$, the permutation must fix the index i, $\sigma'(i) = i$. From the previous part of the proof, we know that $g' \cdot \psi_t(c) = \sigma' \cdot \psi_t(c)$ and since $\sigma'(i) = i$, it follows that $g' \in G_{f'_i}$, where f'_i is the i^{th} fractional coordinate of $\psi_t(c)$. Therefore, $\psi_t(c)$ retains the same site-symmetry structure as c, and is thus also $\mathcal{W}$-constructable.

□

A.3 Proof of Lemma 3.1

Let $g \in G$. Our goal is to show that there exists a permutation $\sigma \in S_n$ such that $g \cdot \boldsymbol{F}_0 = \sigma \cdot \boldsymbol{F}_0$. Since F_0 is $\mathcal{W}$-constructable it can be written as a union of orbits under the action of G. Consider a single orbit generated by w_i, denoted by $\boldsymbol{F}_0^{w_i}$. We observe that $g \cdot \boldsymbol{F}_0^{w_i} = \sigma' \cdot \boldsymbol{F}_0^{w_i}$ for some $\sigma' \in S_{|w_i|}$. This holds because the action of a group element on an orbit is a bijection. Repeating this process for each orbit contained in $\boldsymbol{F}_0$ yields a construction for the permutation σ. □

A.4 Proof of Lemma 3.4

Proof.

$$\begin{aligned}
\hat{u}_t(c|G) &= \frac{1}{|G|} \sum_{g \in G} g \cdot u_t(g^{-1} \cdot c) \\
&= \frac{1}{|G|} \sum_{g \in G} g \cdot u_t(\sigma_{g^{-1}|c} \cdot c) \\
&= \frac{1}{|G|} \sum_{g \in G} g \cdot \sigma_{g^{-1}|c} \cdot u_t(c)
\end{aligned}$$

The second equality follows from the G-symmetry of c, and the final equality follows from the S_n-equivariance of u_t. □

B Conditional Flow on Fractional Coordinates

Lemma B.1. *Let G be a space group, the flat tori logarithmic map $\log_{\boldsymbol{F}_0}(\boldsymbol{F}_1)$ is $G \times S_n$ equivariant.*

Proof. Let $g \in G$ such that $g = (\boldsymbol{R}, \tau)$. Since the orthogonal components of G map the crystal to itself, they preserve the lattice structure. Combining this observation with the following lemma (which is expressed with respect to a single point), we obtain that the logarithmic map is equivariant with respect to the space group and that the representation that acts on the output domain includes only the orthogonal part without the translation. The S_n equivariance is trivial for an element-wise function. □

Lemma B.2. *Let $g \in G$ such that $g = (\boldsymbol{R}, \tau)$. If $\boldsymbol{R}$ maps $\mathbb{Z}^3$ to itself, then $\boldsymbol{R} \log_x(y) = \log_{g \cdot x}(g \cdot y)$.*

Proof. Let $\log_x(y) = v$. This means that there exists $z \in \mathbb{Z}^3$ s.t $v = y - x + z$ where $v \in [-\frac{1}{2}, \frac{1}{2})^3$. Now let's assume $\log_{g \cdot x}(g \cdot y) = v'$. This means that there exists $z' \in \mathbb{Z}^3$ s.t $v' = g \cdot y - g \cdot x + z'$. Plugging in $g = (\boldsymbol{R}, \tau)$ results in $v' = \boldsymbol{R}(y - x) + z'$. Combining both equations, we obtain $v' = \boldsymbol{R}v + z''$ (because $\boldsymbol{R}z \in \mathbb{Z}^3$). Since $v' \in [-\frac{1}{2}, \frac{1}{2})^3$ and $\|v\| = \|\boldsymbol{R}v\|$ we conclude that $v' = \boldsymbol{R}v$.

□

Lemma B.3. *The conditional flow $\psi_t(\boldsymbol{F}_0|\boldsymbol{F}_1)$ is G-symmetric and $\mathcal{W}$-constructable.*

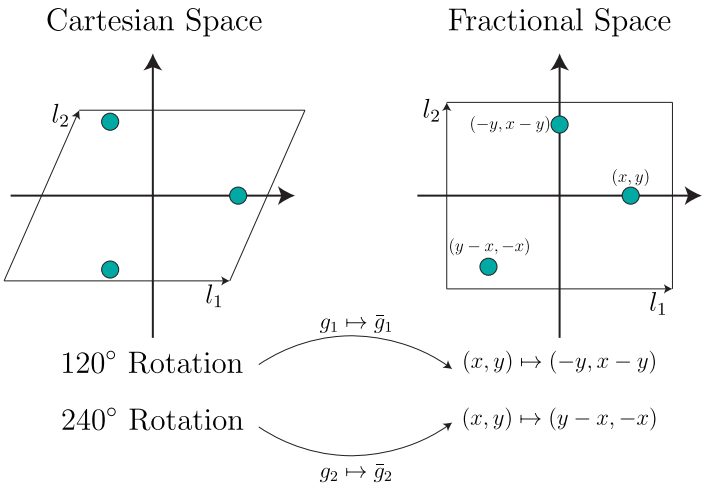


Figure 4: 2D visualization of the deformed geometry induced by moving to fractional coordinates in a non-orthogonal lattice basis. This example demonstrates how a 3-fold rotational space group becomes a set of special affine transformations when acting on fractional coordinates and, specifically, the mapping of each rotation element to a corresponding special affine transformation. In addition, the figure illustrates how a space group induced orbit transforms under this change of basis. This new representation makes it difficult to use models that rely on Euclidean geometry.

Proof.

$$\begin{aligned} g \cdot \psi_t(\boldsymbol{F}_0|\boldsymbol{F}_1) &= (\boldsymbol{F}_0 + t \log_{\boldsymbol{F}_0}(\boldsymbol{F}_1))R^T + \mathbf{1}_n \tau^T \\ &= \boldsymbol{F}_0 \boldsymbol{R}^T + \mathbf{1}_n \tau^T + t \log_{\boldsymbol{F}_0}(\boldsymbol{F}_1) R^T \\ &= g \cdot \boldsymbol{F}_0 + t \log_{g \cdot \boldsymbol{F}_0}(g \cdot \boldsymbol{F}_1) \\ &= \sigma \cdot \boldsymbol{F}_0 + t \log_{\sigma \cdot \boldsymbol{F}_0}(\sigma \cdot \boldsymbol{F}_1) \\ &= \sigma \cdot \psi_t(\boldsymbol{F}_0|\boldsymbol{F}_1) \end{aligned}$$

The fact that $\psi_t(\boldsymbol{F}_0|\boldsymbol{F}_1)$ is $\mathcal{W}$-constructable follows directly from the proof in Appendix A.2 and the fact that $\psi_t(\boldsymbol{F}_0|\boldsymbol{F}_1)$ is mutually G-symmetric with $\boldsymbol{F}_0$ and $\boldsymbol{F}_1$. □

C Lattice Representation

The lattice matrix $\boldsymbol{L} \in \mathbb{R}^{3\times3}$ characterizes the geometry of the unit cell. When $\boldsymbol{L}$ corresponds to a physically valid lattice, i.e., it has positive volume, it is invertible and can be decomposed into the product $\boldsymbol{L} = \boldsymbol{Q}\exp(\boldsymbol{S})$ where $\boldsymbol{Q} \in \mathbb{R}^{3\times3}$ is an orthogonal matrix and $\boldsymbol{S} \in \mathbb{R}^{3\times3}$ is a symmetric matrix. Representing the lattice parameters via $\boldsymbol{S}$ enjoys the benefits of orthogonal invariance (any orthogonal transformation is added to $\boldsymbol{Q}$), which makes this representation invariant to any space group operations. Jiao et al. (2024) suggested representing $\boldsymbol{S}$ using the coefficients of the following basis -

$$\boldsymbol{B}_1 = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \boldsymbol{B}_2 = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix}, \boldsymbol{B}_3 = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix},$$

$$\boldsymbol{B}_4 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \boldsymbol{B}_5 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & -2 \end{pmatrix}, \boldsymbol{B}_6 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

This basis enables clustering of the crystallographic space groups based on the basis coefficients used to represent $\boldsymbol{S}$. Table 3 summarizes the lattice and coefficient constraints for each crystal family type.

Table 3: Relationship between the lattice shape and the constraint of the symmetric bases.

Crystal Family	Space Group No.	Lattice Shape	Constraint of Symmetric Bases
Triclinic	$1 \sim 2$	No Constraint	No Constraint
Monoclinic	$3 \sim 15$	$\alpha = \gamma = 90^\circ$	$k_1 = k_3 = 0$
Orthorhombic	$16 \sim 74$	$\alpha = \beta = \gamma = 90^\circ$	$k_1 = k_2 = k_3 = 0$
Tetragonal	$75 \sim 142$	$\alpha = \beta = \gamma = 90^\circ$ $a = b$	$k_1 = k_2 = k_3 = 0$ $k_4 = 0$
Hexagonal	$143 \sim 194$	$\alpha = \beta = 90^\circ, \gamma = 120^\circ$ $a = b$	$k_2 = k_3 = 0, k_1 = -log(3)/4$ $k_4 = 0$
Cubic	$195 \sim 230$	$\alpha = \beta = \gamma = 90^\circ$ $a = b = c$	$k_1 = k_2 = k_3 = 0$ $k_4 = k_5 = 0$

D Density Functional Theory

Crystals exist in competition for stability between alternatives with the same composition. If one plots energy against composition, the lowest energy structures form a *convex hull*. We say a crystal is thermodynamically stable if it is near or below this convex hull. Since we do not know all structures, there is epistemic uncertainty in this characterization. Ultimately we report an empirical approximation to the convex hull that may contain no points on the "true" convex hull; however, it is the best estimate we know of. The difference between the energy of a crystal and this convex hull is denoted E_{hull}. We report $E_{\text{hull}} < 100$ meV/atom and $E_{\text{hull}} < 0$ meV/atom rates for stability metrics. These values are computed by prerelaxation with a machine learning interatomic potential Barroso-Luque et al. (2024); Wood et al. (2025) followed by relaxation and energy evaluation using density functional theory.

For the stability metrics, we applied the Vienna ab initio simulation package (VASP) Kresse & Furthmüller (1996) to compute relaxed geometries and ground state energies at a temperature of 0 K and pressure of 0 atm. We used the default settings from the Materials Project Jain et al. (2013) known as the `MPRelaxSet` with the PBE functional Perdew et al. (1996) and Hubbard U corrections. These correspond with the settings that our prerelaxation network OMat24 Barroso-Luque et al. (2024) was trained on, so prerelaxation should reduce DFT energy, up to fitting error.

We did *not* make any guesses about oxidation states! This deviates from the Materials Project which does make those guesses. For this reason, our energy above hull calculations for structures that need to consider oxidation state are slightly high, implying that we might be under-predicting stability. This applies to any stability result we calculated. We expect it to also be a negligible effect.

The results from DiffCSP++ WyFormer in Table 1 were computed by Kazeev et al. (2025) and differ slightly from ours. Specifically, they perform multiple relaxations to avoid errors that come from using a poor initial guess before relaxation. Since we prerelax with OMat24 we expect that double relaxation is unnecessary. Consult their work for further details, but we believe the differences are negligible for this purpose.

We contribute a DNG model trained on MP-20. We evaluate against a Materials Project reference hull called `2023-02-07-ppd-mp.pkl.gz` distributed by `matbench-discovery` (Riebesell et al., 2023).

E Related Work

As a continuation from Section 5, we discuss other related work. We still limit the focus to the most relevant parts of this large body of literature.

Our method resembles non-deep learning based methods that propose structures using Wyckoff positions as inductive bias Glass et al. (2006); Pickard & Needs (2011); Meredig & Wolverton (2013); Michel (2014) and refine the atomic positions using density functional theory. This field is known as high-throughput screening of inorganic crystals and it is responsible for generating several important datasets of stable materials Saal et al. (2013); Kirklin et al. (2015); Wang et al. (2021); Schmidt et al.

(2022a,b). Recently, those searches have been sped up by machine learning interatomic potentials that closely approximate density functional theory Merchant et al. (2023).

Now we take a step further away conceptually to discuss methods that are tangentially related to ours. *Crystal-GFN* Mila AI4Science et al. (2023) is a G-flow network that uses the space group, but does not consider Wyckoff positions. *Mat2Seq* Yan et al. (2025) proposes a one-dimensional sequence representation of crystal structures that embeds space group information. *GMTNet* Yan et al. (2024) enforces space group invariance in a crystal property prediction model. Several other works generate crystals without considering multiple types of atom Wirnsberger et al. (2022), or molecule Köhler et al. (2023). Additionally, there is a large and growing body of literature on generative models for materials that do not have general space group equivariance Xie et al. (2021); Yang et al. (2023); Zeni et al. (2023); Miller et al. (2024); Sriram et al. (2024); Lin et al. (2024); Joshi et al. (2025); Hoellmer et al. (2025).

F Evaluation Details

De Novo Generation In this task, we assess the generative model's capacity to produce thermodynamically stable, unique, and novel crystals (S.U.N.). Here, thermodynamic stability implies a structure is at or near a local minima in composition space, see Appendix D for more information; uniqueness is measured across the generated ensemble; and novelty signifies that the structures are absent from both the training and validation sets, computed using `StructureMatcher` (Ong et al., 2013) with default settings (`ltol=0.2, stol=0.3, angle_tol=5`).

We additionally reported S.S.U.N. results. This analysis requires detecting the symmetry of a DFT-relaxed structure. We used `spglib`'s symmetry determination software with parameters `angle_tolerance=5.0` and `symprec=0.1` (Atsushi Togo & Tanaka, 2024).

Additional proxy metrics are also reported as part of the evaluation. *Structural validity* implies that the pairwise atomic distances of a crystal's atoms are all greater than 0.5Å. *Compositional validity* implies that a crystal has a neutral charge according to the SMACT Davies et al. (2019) rules. The *properties* that we consider for computing divergences include: ρ the atomic density defined by number of atoms divided by unit cell volume, elem (arity) the number of unique elements in a crystal, and cn (coordination number) or the number of bonds per atom on average.

Crystal Structure Prediction For evaluation, a crystal structure is generated for each entry in the test set and then compared against the corresponding ground truth structure using pymatgen `StructureMatcher` (Ong et al., 2013) with the same threshold values as in Jiao et al. (2024) `ltol=0.3, stol=0.5, angle_tol=10`. We report two metrics: the match rate (MR), defined as the fraction of generated structures that successfully match their ground truth counterparts, and the RMSE, averaged over all matched pairs.

G Baselines

We provide additional context on the core approach behind each baseline we compared against. *CDVAE* integrates a diffusion model with a variational autoencoder for crystal structure generation. *ADiT* uses a latent-based diffusion model and is trained on additional information from the QM9 Wu et al. (2018) dataset. *FlowMM* is an application of Riemannian Flow Matching Chen & Lipman (2024) that incorporates non-trivial geometries in the crystal representation space. *FlowLLM* combines FlowMM with a large language model that serves as a base-distribution generator. *OMatG* leverages Stochastic Interpolants Albergo et al. (2023) for material generation. *SymmCD* operates on the asymmetric unit and incorporates Wyckoff positions as part of the generative process. *DiffCSP++* is a diffusion-based model that conditions on space groups and projects each denoising step through Wyckoff position transformations. *WyFormer* employs an autoregressive model to generate Wyckoff positions conditioned on the space group and subsequently uses DiffCSP++ for full crystal-structure generation. Finally, *SGEquiDiff* is a diffusion-based model that enforces space-group equivariance while operating on the asymmetric unit.

H Model Details

H.1 Architecture

In this section, we present a comprehensive overview of our vector field model $\hat{u}_t(\cdot \mid G)$, along with the hyperparameters employed during training and generation across all experiments. The model takes as input a crystal $c = (k, \boldsymbol{F}, \boldsymbol{A})$, where $f^i \in \mathbb{R}^3$ denotes the i^{th} fractional coordinate in $\boldsymbol{F}$, and $a^i \in \{0,1\}^h$ represents the i^{th} atom type indicator vector in $\boldsymbol{A}$. The forward computation of the s-layer model $\hat{u}(c \mid G)$ is defined by the following set of equations:

$$
\begin{aligned}
a^i_{embed} &= \phi^a(a^i) \\
t_{embed} &= \text{SinusoidalTimeEmbedding}(t) \\
h^i_{(0)} &= \phi^{\text{embed}}(\left[a^i_{embed}, t_{embed}\right]) \\
m^{ij}_{(l)} &= \phi^{\text{edge}}_{(l)}(\left[h^i_{(l-1)}, h^j_{(l-1)}, k, \text{SinusoidalEmbedding}(\log_{f^i}(f^j)), \frac{L^T L \log_{f^i}(f^j)}{\left\|L^T L \log_{f^i}(f^j)\right\|}\right]) \\
m^i_{(l)} &= \frac{1}{n}\sum_{j=1}^{n} m^{ij}_{(l)} \\
h^i_{(l)} &= \phi^{\text{node}}_{(l)}(\left[h^i_{(l-1)}, m^i_{(l)}\right]) \\
u^k_t(c_t) &= \phi^k(\frac{1}{n}\sum_{j=1}^{n} h^j_{(s)}) \\
(u^{\boldsymbol{F}}_t(c_t))^i &= \phi^{\boldsymbol{F}}(h^i_{(s)}) \\
(u^{\boldsymbol{A}}_t(c_t))^i &= \phi^{\boldsymbol{A}}(h^i_{(s)}) \\
u_t(c_t) &= (u^k_t(c_t), u^{\boldsymbol{F}}_t(c_t), u^{\boldsymbol{A}}_t(c_t)) \\
\hat{u}_t(c_t \mid G) &= \frac{1}{|G|}\sum_{g \in G} g \cdot \sigma_{g^{-1}|c} \cdot u_t(c_t)
\end{aligned}
$$

We denote d as the hidden dimension of the model, d_t as the Sinusoidal Time Embedding dimension, and d_s as the Sinusoidal Embedding dimension. Next, we list the learnable modules constructing the model and denote their input and output dimension as $x \to y$. ϕ^a is a linear layer $h \to d$, ϕ^{embed} is a linear layer $d + d_t \to d$, $\phi^{\text{edge}}_{(l)}$ is 2-layer Multi-layer Perceptron (MLP) $2d + d_s + 9 \to d$, $\phi^{\text{node}}_{(l)}$ is a 2-layer Multi-layer Perceptron (MLP) $2d \to d$, ϕ^k is a linear layer $d \to 6$, $\phi^{\boldsymbol{F}}$ is a linear layer $d \to 3$ and $\phi^{\boldsymbol{A}}$ is a linear layer $d \to h$. The last equation represents the group averaging presented in Equation (4). The flat tori logarithmic map is defined by the equation:

$$
\log_{f^i}(f^j) = \frac{1}{2\pi}\,\text{atan2}(\left[\sin(f^j - f^i), \cos(f^j - f^i)\right]) \tag{11}
$$

Table 4 summarizes the hyperparameters used to train our SGFM models. Note that the same configuration was applied uniformly across all datasets and tasks. The hyperparameter search was performed on MP-20 (CSP) and the resulting settings were adopted for all other experiments.

H.2 Training & Generation

All of our models were trained using the `ADAM` optimizer Kingma & Ba (2014) on 8 NVIDIA A100 GPUs. Table 5 outlines the training configuration for each model, including the ranges explored during hyperparameter search. We employed a cosine annealing learning rate schedule with a minimum learning rate of 0.00001. As described in Section 4, we applied inference anti-annealing to enhance generation quality. This technique modifies the vector field by scaling it with a time-dependent function $s(t) = 1 + s't$, where $s' \in \mathbb{R}^+$ is a hyperparameter. We defined separate annealing parameters for each crystal component: $s'_{\boldsymbol{F}}$ and s'_k (no annealing was applied to atom type prediction). All datasets have $60/20/20$ train/validation/test split except Alex-MP-20 that has $80/10/10$ split. We used the same split exact indexing split as OMatG to produce a train, test, and validation set Hoellmer et al. (2025). OMatG derived its train/test split from MatterGen Zeni et al. (2023) where it took 10% of the training data to make a validation set.

Table 4: Hyperparameter details for all the models reported in the paper. Hyperparameter (bottom row) search was conducted on the MP-20 dataset.

Dataset	Number of Layers	d	d_t	d_s	Activation	Layer Norm
			CSP			
MP-20	8	512	256	128	SiLU	✓
MPTS-52	8	512	256	128	SiLU	✓
Carbon-24	8	512	256	128	SiLU	✓
Perov-5	8	512	256	128	SiLU	✓
Alex-MP-20	8	512	256	128	SiLU	✓
			DNG			
MP-20	8	512	256	128	SiLU	✓
			Hyperparameter Range			
MP-20	$\{6, 7, 8, 9\}$	$\{128, 256, 512\}$	$\{128, 256\}$	$\{128, 256\}$	-	-

Table 5: Training hyperparameter details for all the models reported in the paper. Hyperparameter (bottom row) search was conducted per experiment.

Dataset	Batch Size/GPU	Learning Rate	Epochs	$\lambda_{\boldsymbol{F}}$	$\lambda_{\boldsymbol{A}}$	λ_k	$s'_{\boldsymbol{F}}$	s'_k
			CSP					
MP-20	64	0.0005	5000	100	-	1	3	3
MPTS-52	32	0.0005	5000	100	-	1	3	3
Carbon-24	64	0.0005	8000	100	-	1	3	3
Perov-5	256	0.0005	1000	100	-	1	3	3
Alex-MP-20	32	0.0003	1250	100	-	1	5	3
			DNG					
MP-20	64	0.0007	5000	100	1	1	5	3
			Hyperparameter Range					
MP-20	-	$\{0.0002, 0.0005, 0.0007\}$	$\{1000, 1250, 2000, 5000, 8000\}$	$\{1, 10, 50, 100\}$	$\{1, 10, 50, 100\}$	$\{1, 10, 50, 100\}$	$\{0, 1, 3, 5\}$	$\{0, 1, 3, 5\}$

H.3 Atom Type Prediction Properties

This section explains the properties of the atom-prediction component of $\hat{u}(\cdot \mid G)$. Specifically, it shows that when the input is a symmetric crystal c, the predicted atom-type vector field $\hat{u}^{\mathcal{A}}(c \mid G)$ assigns identical values to atoms belonging to the same orbit. A point that may not be immediately obvious is that the group actions are defined in a similar way on the output $\hat{u}(c|G) = (\hat{u}^k(c|G), \hat{u}^{\boldsymbol{F}}(c|G), \hat{u}^{\boldsymbol{A}}(c|G))$ as they are defined over c, up to a representation, as discussed in the section on fractional coordinates.

To begin with, we can show that $\hat{u}^{\mathcal{A}}(c \mid G)$ is G–invariant. On one hand, we know that $\hat{u}^{A}(c \mid G) = \hat{u}^{A}(g \cdot c \mid G)$ because $\hat{u}$ is G–equivariant and the group element g does not act directly on the atom-type component. Moreover, we also have $\hat{u}^{A}(g \cdot c \mid G) = \hat{u}^{A}(\sigma_{g|c} \cdot c \mid G) = \sigma_{g|c} \cdot \hat{u}^{A}(c \mid G)$ since $\hat{u}$ is also S_n–equivariant and c is G–symmetric.

This relation, $\hat{u}^{A}(c \mid G) = \sigma_{g|c} \cdot \hat{u}^{A}(c \mid G)$, applies only to permutations corresponding to elements of the space group, which only permute atoms within the same orbit induced by the group action. Consequently, we can conclude that the atom-type vector field $\hat{u}_t^{\mathcal{A}}$ takes identical values for atoms belonging to the same orbit.

I Ablations

I.1 Running Time Ablations

Table 6: Generation times for various batch sizes.

Model	Generation Time (s/batch)		
	64	256	500
DiffCSP++	66	268	490
SGEquiDiff	19	75	154
SGFM	22	92	175

Table 7: Training and generation time comparison of different vector field models.

Model	Batch size	Training Time (s)	Generation Time (s)
SGFM	64	28.2	17.81
Non-Equivariant	64	26.3	16.39
GA	1	600	14.2

Table 8: **CSPML evaluation.** Wyckoff positions are obtained via the CSPML heuristic. The reported metrics are MR and RMSE.

Model	MP-20		MPTS-52		Perov-5	
	MR (%) ↑	RMSE ↓	MR (%) ↑	RMSE ↓	MR (%) ↑	RMSE ↓
CSPML	70.51	.0388	**36.98**	.0664	51.84	.1066
DiffCSP++	**70.58**	.0272	37.17	.0676	52.17	.0841
SGFM	70.13	.0310	35.09	.0739	**54.10**	.1508

This experimental ablation study aims to evaluate the efficiency of SGFM by comparing its performance during both training and generation against two baseline models: (1) a non-equivariant variant where the vector field does not incorporate group symmetry, and (2) a standard GA implementation as defined in Equation (3). For each model, we measured the time required to train a single epoch on MP-20, as well as the time needed to generate a batch of 64 (with 500 generation steps). The training time was averaged over 10 epochs, while the generation time was averaged over 100 batches. Training was conducted on an NVIDIA RTX8000 using 8 GPUs, while generation was performed on a single NVIDIA A10 GPU. The results are summarized in Table 7. Due to memory constraints, the standard GA model was limited to a maximal batch size of 1 per GPU. The generation results for the standard GA model were obtained from a randomly initialized model, as training such a model was intractable. The results highlight a significant efficiency gap between the SGFM implementation of GA and the standard version, while showing minimal difference compared to the non-equivariant model.

Table 9: **Noise prior vs. model equivariance comparison.** The reported result is MR on the MP-20 dataset.

Noise Prior \Model	Equivariant	Non-Equivariant
Wyckoff Conditional	**82.74**	68.16
Uniform	-	64.49

I.2 Additional Results

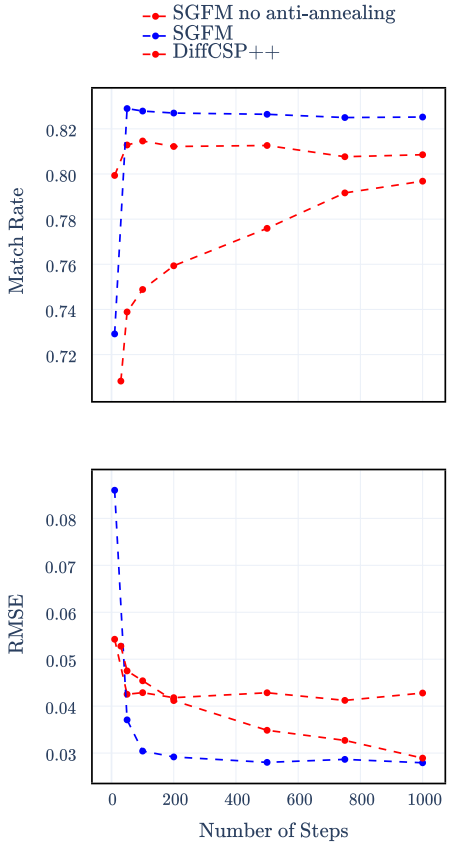


Figure 5: **Sampling Efficiency Analysis.** MR (↑) and RMSE (↓) versus integration steps on MP-20. We observe that SGFM reaches near-optimal accuracy within just 50–100 steps. DiffCSP++ converges more slowly, requiring up to 1000 steps to approach its best performance—while still showing a notable gap in match rate compared to SGFM. In this same plot, we include a version of SGFM without anti-annealing to assess the effects of anti-annealing on performance.

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: [N/A]

   Guidelines:

   - The answer [N/A] 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 [N/A] 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: The conclusion section addresses the limitations of our method and suggests future directions to improve it.

   Guidelines:

   - The answer [N/A] 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: [N/A]

Guidelines:

- The answer [N/A] 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: See Appendix H

   Guidelines:

   - The answer [N/A] 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: [No]

Justification: [N/A]

Guidelines:

- The answer [N/A] means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (`https://neurips.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://neurips.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) necessary to understand the results?

   Answer: [Yes]

   Justification: See Appendix F

   Guidelines:

   - The answer [N/A] 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: The results in the paper are reported over an independent test set following the evaluation protocol of previous work.

   Guidelines:

   - The answer [N/A] 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: See Appendix F

   Guidelines:

   - The answer [N/A] 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:

   Guidelines:

   - The answer [N/A] 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: [N/A]

    Justification:

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.
    - If the authors answer [N/A] 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., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification:

    Guidelines:

    - The answer [N/A] 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: [N/A]

    Justification:

    Guidelines:

    - The answer [N/A] 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: [N/A]

    Justification:

    Guidelines:

    - The answer [N/A] 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: [N/A]

    Justification:

    Guidelines:

    - The answer [N/A] 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: [N/A]

    Guidelines:

    - The answer [N/A] 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 rigor, or originality of the research, declaration is not required.

Answer: [N/A]

Justification:

Guidelines:

- The answer [N/A] 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 in the NeurIPS handbook for what should or should not be described.