DYNACRYS: CRYSTAL GENERATION WITH DYNAMIC SPACE-GROUP DIFFUSION

Anonymous authors
Paper under double-blind review

ABSTRACT

The search for new crystalline materials spans an enormous compositional and structural space. Generating candidates in this space requires modeling the relationships between elemental composition, crystallographic symmetry, and crystal geometry. The space group specifies the crystal's symmetry and determines the admissible Wyckoff positions and geometric degrees of freedom; fixing it in advance prevents this global symmetry choice from being revised using information from evolving Wyckoff occupations and elemental composition. We introduce DynaCrys, the first framework for complete crystal generation in which the space group co-evolves with Wyckoff occupations and elements. We develop physics-informed space-group diffusion that explicitly incorporates crystallographic group–subgroup relations, and a legality-constrained stochastic decoder that jointly samples compatible space groups, Wyckoff occupations, and elements, followed by symmetry-constrained geometry generation. In large-scale evaluations on standard benchmarks, DynaCrys achieves state-of-the-art performance in symmetry-aware crystal generation while supporting fast sampling.

1 INTRODUCTION

Discovering new crystalline materials requires searching a vast space of elemental compositions and atomic arrangements. Generative models offer a way to propose candidates directly, but the variables that define a crystal are tightly coupled. Its space group G determines the admissible Wyckoff positions and their geometric degrees of freedom. The space group, Wyckoff occupations, and their elements define the crystal's protostructure—its symbolic framework. Lattice parameters and atomic coordinates specify its geometry within the corresponding symmetry constraints.

In crystal generation, models such as MatterGen (Zeni et al., 2025) and FlowMM (Miller et al., 2024) directly model lattice parameters, atomic coordinates, and elements, with space groups determined post hoc from the generated structures. In contrast, symmetry-aware models explicitly use space groups, Wyckoff occupations, or site symmetries to constrain generation (Table 1). These models, however, select and fix the space group before generating the remaining variables, restricting subsequent exploration to configurations compatible with that group. Joint generation can instead use information from the evolving occupations and elemental composition to update the space-group choice.

Table 1: Crystallographic symmetry across models.

Methods	Symmetry-aware	Dynamic G	Group–subgroup structure
MatterGen, FlowMM, DiffCSP	×	×	×
DiffCSP++, SymmCD, WyFormer SGEquiDiff, WyckoffDiff	✓	×	×
DynaCrys (ours)	✓	✓	✓

We introduce DynaCrys, a framework for complete crystal generation in which the space group jointly evolves and is dynamically revised with Wyckoff occupations and elements through diffusion. The physics-informed space-group diffusion kernel incorporates crystallographic group–subgroup relations (Bärnighausen, 1980), which describe changes in symmetry constraints and their associated degrees of freedom. A legality-constrained stochastic decoder then samples compatible symbolic

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

states. Conditioned on the final state, a symmetry-constrained geometry model generates lattice parameters and atomic coordinates to produce a complete crystal (Figure 1).

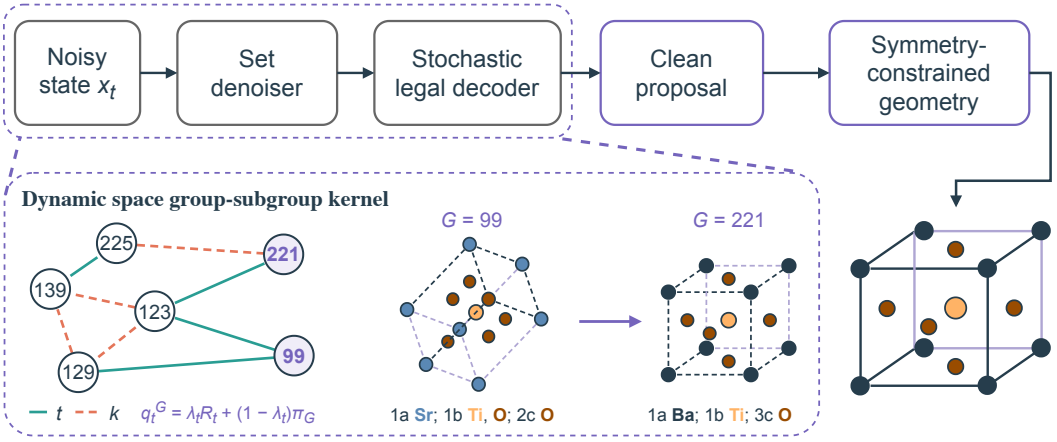


Figure 1: Overview of DynaCrys: joint generation of the space group, Wyckoff occupations, and elements, followed by symmetry-constrained geometry generation.

Our contributions are:

- *Joint crystal generation with a dynamic space group.* We introduce the first complete crystal generator to jointly evolve the space group, Wyckoff occupations, and elements during generation.
- *Physics-informed space-group diffusion.* We incorporate the group-theoretic structure of crystal symmetries into the space-group diffusion kernel, using crystallographic group–subgroup relations as a structured prior for joint generation.
- *Legality-constrained stochastic decoding.* We develop a joint stochastic decoder. It can sample according to model-predicted probabilities conditioned on space-group and occupation constraints, without costly summation over all legal configurations.
- *Material discovery and validation.* We demonstrate state-of-the-art generation performance among symmetry-aware crystal generative models on MP-20 and MPTS-52. Independent evaluations using multiple machine-learning interatomic potentials and large-scale first-principles calculations consistently confirm that DynaCrys generates a higher proportion of stable, unique, and novel crystals.

2 Related Work

Crystal generation without explicit space-group modeling. CDVAE (Xie et al., 2022), DiffCSP (Jiao et al., 2023), and MatterGen (Zeni et al., 2025) use diffusion-based generative models, while FlowMM (Miller et al., 2024) uses Riemannian flow matching. These methods directly model lattice parameters, atomic coordinates, and elements. They use periodic representations and geometric invariance or equivariance to handle global translations and rotations. Their space groups are identified post hoc from the generated structures rather than represented as explicit generative variables.

Crystal generation with explicit crystallographic symmetry. Explicitly using space groups, Wyckoff positions, or site symmetries allows generative models to enforce crystallographic constraints directly. Existing methods determine orbit occupations through different mechanisms. Template-based DiffCSP++ (Jiao et al., 2024) directly obtains space groups and occupation patterns from training crystal data, then assigns elements to symmetry orbits and generates geometry. Other methods learn to generate occupations: SymmCD (Levy et al., 2025) determines Wyckoff positions by generating site symmetries; WyFormer (Kazeev et al., 2025) and SGEquiDiff (Chang et al., 2025) generate occupations and elements autoregressively; and WyckoffDiff (Ekström Kelvinius et al., 2025) diffuses discrete occupation variables conditioned on the space group. These methods select and fix the space group before generating the remaining variables. DynaCrys jointly updates the space

group, Wyckoff occupations, and elements through diffusion guided by group–subgroup relations, allowing the space-group choice to change during generation (Table 1).

Constrained decoding for crystal generation. Existing crystal generative models handle constraints in different ways. SGEquiDiff (Chang et al., 2025) prevents repeated occupation of 0D Wyckoff positions during autoregressive sampling; WyckoffDiff (Ekström Kelvinius et al., 2025) enforces this restriction through its occupation representation; and SymmCD (Levy et al., 2025) obtains valid site symmetries and coordinates through projection. DynaCrys uses legality-constrained joint stochastic decoding to sample compatible states because it needs to coordinate the dynamic space group and other variables, ensuring that the predicted orbit count matches the number of selected orbit occupations and that occupation-capacity and atom-count constraints are satisfied.

3 Preliminaries: Space Groups and Wyckoff Orbits

A periodic crystal consists of a lattice basis $L \in \mathbb{R}^{3\times3}$ and atoms in one unit cell, represented here by fractional coordinates $\mathbf{f} \in [0,1)^3$. Rather than treating these atoms as independent points, a space group G organizes them into symmetry orbits. There are 230 three-dimensional space-group types (Aroyo, 2016). Each operation $g = (R_g, \mathbf{t}_g)$ acts on fractional coordinates as $g \cdot \mathbf{f} = (R_g\mathbf{f} + \mathbf{t}_g) \bmod 1$. The orbit of a representative coordinate is the set of distinct symmetry-equivalent sites obtained by applying all operations in G.

A Wyckoff position describes a family of symmetry orbits. Each Wyckoff position corresponds to a row in the space-group table, indexed by $r \in \mathcal{R}_G$. The row specifies the orbit multiplicity $\mu_{G,r}$ (the number of equivalent sites in an orbit) and the number of free coordinates $\mathrm{dof}_{G,r}$. Choosing values for the free coordinates selects a particular orbit; a row with no free coordinates already determines the orbit. For example, r can denote the row labeled $4e$ in Figure 2 (left). The four matching diamonds form one orbit of this Wyckoff position. Here, 4 gives the multiplicity and e identifies the Wyckoff position within G (Aroyo, 2016). Varying the free coordinate in this example can produce another orbit represented by the same row.

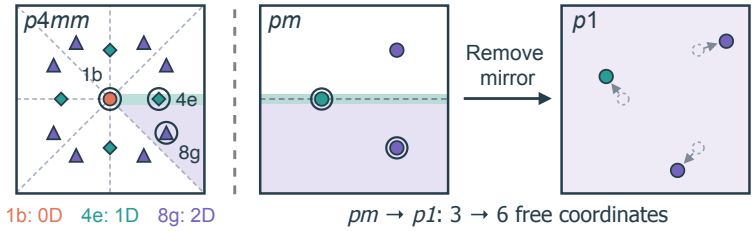


Figure 2: Schematic of Wyckoff orbits and symmetry constraints. Left: $p4mm$ orbits with circled representatives; matching markers denote equivalent sites (inspired by Jiao et al., 2024). Right: orbit splitting under $pm \to p1$, with dashed circles marking the original positions.

A symbolic orbit occupation specifies a row r and an element Z, leaving the free-coordinate values unspecified. Each occupation therefore contributes $\mu_{G,r}$ atoms of element Z. Rows with free coordinates may be occupied more than once. Together with G, these occupations define the protostructure; geometry generation supplies the lattice parameters and remaining free coordinates.

The vocabulary $\mathcal{R}_G$ is specific to the space group. Wyckoff labels are defined within each space group. Changing G changes the available rows, their multiplicities, and their free-coordinate parameterizations. Group–subgroup relations organize these changes. A maximal relation connects two space groups with no intermediate space group. Every such step is either translationengleiche (t), which preserves the translation lattice while reducing point group symmetry, or klassengleiche (k), which retains the point group while reducing translational symmetry (Aroyo, 2016). Restricting G to a subgroup H can split one G-orbit into several independently parameterized H-orbits. In the illustrated two-dimensional $pm \to p1$ transition (Figure 2, right), the three atoms can then vary independently, increasing the number of free coordinates from 3 to 6.

Consequently, revising the space group during symbolic generation does more than update a categorical label: it changes the admissible orbit representation and the continuous degrees of freedom available to the subsequent geometry stage. Section 4 therefore treats G, the orbit occupations, and their elements as coupled components of one evolving symbolic state.

4 DYNAMIC SPACE-GROUP CRYSTAL GENERATION

DynaCrys separates crystal generation into two stages. A symbolic diffusion model first samples a protostructure containing the space group, Wyckoff occupations, and their elements. These variables evolve jointly during this stage. Conditioned on this symbolic state, a symmetry-constrained geometry model then generates the remaining lattice parameters and free coordinates.

4.1 PROTOSTRUCTURE STATE AND SHARED SYMMETRY REPRESENTATION

Let $\mathcal{G}$ denote the space-group vocabulary. A protostructure P (Goodall et al., 2022) is

$$P = \big(G,\ K,\ \{(Z_i, r_i)\}_{i=1}^{K}\big), \tag{1}$$

where $G \in \mathcal{G}$ is the space group, K is the number of orbit occupations, $r_i \in \mathcal{R}_G$ is a Wyckoff row, and Z_i is the element assigned to that occupation. The occupations form an unordered multiset, the same row may occur more than once, and P contains neither a lattice nor free coordinates. The symbolic diffusion represents P with a fixed-width padded clean state x_0. At noise level t, the state is

$$x_t = \big(G_t,\ K_t,\ \{(m_{i,t}, Z_{i,t}, y_{i,t})\}_{i=1}^{M}\big). \tag{2}$$

Here $m_{i,t}$ marks an active or padded slot, $Z_{i,t}$ is the corresponding element or padding symbol, and $y_{i,t} \in \mathbb{R}^d$ is a noisy Wyckoff-row vector. At $t = 0$, the active slots encode the occupations of P and the remaining slots are padding. The noisy state x_t need not decode directly to a legal protostructure.

Because the row vocabulary $\mathcal{R}_G$ depends on G, DynaCrys uses a shared symmetry codebook containing space-group vectors c_G and Wyckoff-row vectors $c_{G,r}$. The codebook is pretrained from crystallographic descriptors. For an active clean slot, $y_{i,0} = c_{G,r_i}$, linking its discrete row to the continuous row channel. The row vectors also serve as the tied vectors of the space-group-conditional row head. The geometry model is also conditioned on the space-group and row codebook embeddings. At each diffusion step, c_{G_t} represents the current noisy space group G_t. The shared codebook therefore provides a common representation of space-group-specific Wyckoff vocabularies across the symbolic and geometry stages. Appendix C provides details of the codebook construction.

4.2 JOINT SYMBOLIC DIFFUSION

Forward corruption. The symbolic forward process is a diffusion process defined by a family of marginal kernels $q_t(x_t \mid x_0)$ (Austin et al., 2021). Its space-group channel mixes a structured transition with a reset to the training distribution:

$$q_t^G(G_t \mid G_0) = \lambda_t R_t(G_t \mid G_0) + (1 - \lambda_t)\pi_G(G_t). \tag{3}$$

The transition R_t diffuses over the undirected graph induced by maximal translationengleiche and klassengleiche group–subgroup relations on $\mathcal{G}$. This graph supplies the transition geometry of the space-group corruption; the reset component makes the terminal distribution independent of G_0.

After drawing G_t, the forward process samples K_t and partitions the M slots into survivors, births, and padding. Survivors retain a corrupted signal from distinct clean occupations, whereas birth rows are drawn from the vocabulary of the current G_t; the row vectors and elements are then corrupted conditionally on this partition. This construction accommodates a changing number of occupations and defines the same marginal kernel for training-time corruption and sampling-time re-noising. Further details on the channel distributions and noise schedules are provided in Appendix B.

Set denoising network. The denoiser represents x_t by one global token for (G_t, K_t, t) and M orbit-slot tokens containing $(m_{i,t}, Z_{i,t}, y_{i,t})$. A Transformer encoder (Vaswani et al., 2017) processes these tokens without slot positional encodings, making its slot outputs permutation equivariant. Global and slot-wise heads predict the clean state as

$$p_\theta(G_0 \mid x_t), \qquad p_\theta(K_0 \mid x_t), \qquad \big\{p_i^{\mathrm{act}}(\cdot \mid x_t),\ p_i(r \mid G, x_t),\ p_i(Z \mid G, r, x_t)\big\}_{i=1}^{M}. \tag{4}$$

Here $p_\theta(G_0 \mid x_t)$ and $p_\theta(K_0 \mid x_t)$ are the global predictions of the clean space group and orbit count. For slot i, p_i^{act} is the binary distribution over occupied (1) and padded (0) states. The row head $p_i(r \mid G, x_t)$ assigns probability only to valid rows $r \in \mathcal{R}_G$ and uses their shared-codebook vectors as classifier weights. The element head $p_i(Z \mid G, r, x_t)$ conditions its predictions on the selected (G, r) through feature-wise linear modulation (FiLM) (Perez et al., 2018) (Appendix D.2).

Training objective. The symbolic model predicts x_0 with the set objective

$$\mathcal{L}_{\text{sym}} = \mathcal{L}_G + w_K\mathcal{L}_K + w_m\mathcal{L}_{\text{mask}} + w_r\mathcal{L}_r + w_Z\mathcal{L}_Z + s_{\text{aux}}(w_{\text{cap}}\mathcal{L}_{\text{cap}} + w_N\mathcal{L}_N + w_{\text{comp}}\mathcal{L}_{\text{comp}}). \quad (5)$$

Here $\mathcal{L}_G$ and $\mathcal{L}_K$ are cross-entropies for the global space-group and orbit-count predictions. A Hungarian assignment (Carion et al., 2020) matches the M predicted slots to the K_0 clean occupations and $M - K_0$ padding targets before the slot losses are evaluated. After matching, $\mathcal{L}_{\text{mask}}$ is the active-state binary cross-entropy over all slots, while $\mathcal{L}_r$ and $\mathcal{L}_Z$ are the row and element negative log-likelihoods over occupied matches. Among the auxiliary terms, $\mathcal{L}_{\text{cap}}$ penalizes excess expected row occupancy, while $\mathcal{L}_N$ and $\mathcal{L}_{\text{comp}}$ match the expected conventional-cell atom count and composition to the clean target. The sampler separately enforces row-capacity limits and the fixed atom-count cap. A separate fine-tuning stage refines the row-conditioned element head with regularization. Appendix D gives the loss weights, exact assignment, auxiliary-loss definitions, and optimization settings.

4.3 Legality-constrained sampling

The heads in Eq. (4) define a factorized clean-state distribution, but the independently predicted orbit count and active states need not agree, and the selected rows need not jointly satisfy the row-capacity and atom-count constraints. Let $\mathcal{X}_{\text{legal}}$ denote the corresponding legal set: G must lie in $\mathcal{G}$, exactly K slots must be active, every selected row must belong to $\mathcal{R}_G$, and the collection of occupations must satisfy the row-capacity and atom-count constraints. The stochastic decoder targets the factorized distribution conditioned on this set,

$$\widetilde{p}_\theta(x_0 \mid x_t) \propto p_\theta^{\text{fac}}(x_0 \mid x_t)\mathbf{1}[x_0 \in \mathcal{X}_{\text{legal}}]. \quad (6)$$

Here p_θ^{fac} is the factorized distribution induced by the heads in Eq. (4), $\mathbf{1}[\cdot]$ is the legality indicator, and $\widetilde{p}_\theta$ is the resulting distribution restricted to and renormalized over $\mathcal{X}_{\text{legal}}$.

The global constrained normalizer for Eq. (6) is the total probability mass that the factorized model assigns to $\mathcal{X}_{\text{legal}}$. Evaluating it directly would require summing over a combinatorial number of legal symbolic states. Instead, let a_i be the active logit of slot i, let $\omega_i = \exp(a_i)$, and define the elementary symmetric polynomial $e_K(\boldsymbol{\omega}) = \sum_{|A|=K} \prod_{i\in A} \omega_i$, where the sum ranges over $A \subseteq \{1, \dots, M\}$. After drawing G, the decoder samples the orbit count K and active-slot set A according to

$$q_{\text{prop}}(K \mid x_t) \propto p_\theta(K_0 = K \mid x_t) e_K(\boldsymbol{\omega}), \quad q_{\text{prop}}(A \mid K, x_t) = \Big(\prod\nolimits_{i\in A} \omega_i\Big)/e_K(\boldsymbol{\omega}). \quad (7)$$

Here $|A| = K$. The e_K factors cancel when the two proposal distributions are multiplied. After sampling G, K, the active set, and the selected rows, the decoder rejects proposals that violate these constraints and draws their elements after acceptance.

Proposition 1. Conditional on acceptance, the stochastic decoder returns an exact draw from $\widetilde{p}_\theta(\cdot \mid x_t)$ without evaluating its global normalizer.

Appendix E gives the derivation and sampling procedure. Reverse generation repeatedly predicts a legal clean state and then applies the same marginal forward kernel at the next noise level $t' < t$:

$$\hat{x}_0^{(t)} \sim \widetilde{p}_\theta(\cdot \mid x_t), \qquad x_{t'} \sim q_{t'}(\cdot \mid \hat{x}_0^{(t)}). \quad (8)$$

At the final sampled noise level, removing inactive slots from the clean prediction and disregarding slot order yields the protostructure P. The space group is thus resampled jointly with the other symbolic variables throughout the reverse trajectory rather than fixed in advance.

4.4 Symmetry-constrained geometry generation

The final protostructure P specifies the space group, Wyckoff occupations, and elements. The geometry stage generates a lattice compatible with the crystal family of G and the free fractional coordinates of one representative point, or orbit anchor, per occupied orbit.

Geometry denoising. Building on the continuous diffusion framework of DiffCSP++ (Jiao et al., 2024), an atom-level message-passing network operates on the full conventional cell reconstructed

from the occupied orbit anchors and is conditioned on P, the diffusion time, and the shared symmetry codebook. Let b_t denote the noisy lattice representation and X_t the noisy full-cell fractional coordinates. The two heads minimize

$$\begin{aligned}\mathcal{L}_{\text{lat}} &= \mathbb{E}\Big[\|\epsilon_G - \widehat{\epsilon}_\theta(b_t, X_t, P, t)\|_2^2\Big], \qquad \epsilon_G = \Pi_G^{\text{lin}}\epsilon, \\ \mathcal{L}_{\text{coord}} &= \mathbb{E}\Big[\big\|\Pi_{\mathcal{A}(P)}\widehat{u}_\theta(X_t, b_t, P, t) - u^*_{\text{wn}}(X_t, X_0; \sigma_t)\big\|_2^2\Big].\end{aligned} \tag{9}$$

Here $\epsilon \sim \mathcal{N}(0, I)$ is standard Gaussian lattice noise, $\epsilon_G = \Pi_G^{\text{lin}}\epsilon$ is its projection onto the crystal-family subspace, and $\widehat{\epsilon}_\theta$ is the lattice-noise prediction. The clean full-cell fractional coordinates are denoted by X_0, σ_t is the coordinate-noise scale, and u^*_{wn} is the wrapped-normal denoising target in the orbit-anchor coordinates. The coordinate head predicts the denoising field $\widehat{u}_\theta$. Finally, $\Pi_{\mathcal{A}(P)}$ maps each predicted atomic field to the free orbit-anchor coordinates $\mathcal{A}(P)$ specified by P.

Symmetry-constrained sampling. At each reverse step, the lattice representation is projected to the subspace compatible with the crystal family of G. The anchor coordinates are updated only along their allowed degrees of freedom and expanded through the corresponding Wyckoff operations into complete atomic orbits. By construction, every intermediate geometry is invariant under the operations of the decoded space group G. Further details on the lattice representation, geometry objectives, and reverse sampling procedure are provided in Appendix F.

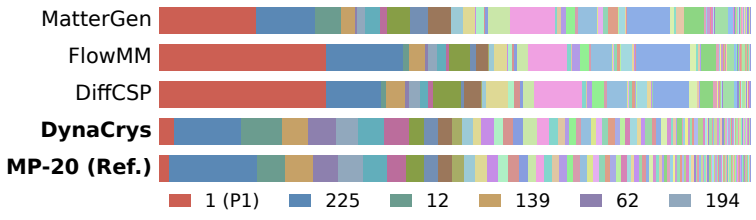
5 Experiments

We evaluate unconditional crystal generation on MP-20 (Xie et al., 2022) and the more challenging MPTS-52 benchmark (Chang et al., 2025). Following standard evaluation practice (Levy et al., 2025), we use each dataset's training split as the novelty reference and its test split as the coverage reference. We evaluate approximately 10k candidates per set using common relaxation, stability, novelty, and symmetry-analysis settings for comparability. The large sample sizes reduce finite-sample variation. We use publicly available generated structures for the baselines. For SGEquiDiff, these were unavailable, so we sample from its released checkpoint. Unless otherwise stated, we report five-seed means for discovery rates and structural statistics of DynaCrys models and SGEquiDiff, use CHGNet (Deng et al., 2023) for relaxation, and identify space groups with `symprec`$= 0.1$ Å and an angular tolerance of $5°$. We also independently evaluate with MACE-MP-0 (Batatia et al., 2025) (Appendix H) and perform large-scale first-principles calculations with VASP to test the robustness of DynaCrys's performance. Full evaluation settings are in Appendix G.

We compare against strong baselines from peer-reviewed publications, including DiffCSP (Jiao et al., 2023), MatterGen (Zeni et al., 2025), FlowMM (Miller et al., 2024), DiffCSP++ (Jiao et al., 2024), SymmCD (Levy et al., 2025), WyFormer (Kazeev et al., 2025), and SGEquiDiff (Chang et al., 2025).

5.1 Symmetry and space-group distributions

We first assess the symmetry diversity of generated crystals. Figure 3 compares the space-group distributions of crystals generated by different models with the MP-20 reference distribution and also reports the fraction of $P1$ structures after relaxation. $P1$ has only translational symmetry and is rare among naturally occurring crystals (Levy et al., 2025).



Method	P1 (%)
MatterGen	20.28
FlowMM	34.02
DiffCSP	30.79
DynaCrys (ours)	5.54

Figure 3: Left: space-group distributions of generated crystals on MP-20. Right: fractions of relaxed structures identified as $P1$.

Figure 3 shows that DynaCrys matches the MP-20 space-group distribution substantially more closely than the symmetry-unaware baselines. These baselines overproduce $P1$ structures, whereas DynaCrys better captures the relative frequencies of a broad range of nontrivial space groups. Its $P1$ fraction also remains substantially lower after relaxation. Together, these results demonstrate a clear advantage over symmetry-unaware methods in matching the space-group statistics of known crystals and generating nontrivially symmetric structures after relaxation.

5.2 Materials discovery

The stable, unique, and novel (SUN) rate (Zeni et al., 2025) is a key metric for evaluating crystal generation. SUN@τ is the fraction of generated candidates that satisfy the stability criterion $E_{\text{hull}} \leq \tau$ ($\tau = 0$ and 0.1 eV/atom), are unique within the generated set, and are novel relative to the training set. SSUN@τ additionally requires a non-$P1$ space group after relaxation (Kazeev et al., 2025).

We report results for DynaCrys and DynaCrys-II, two models that share the same framework and differ only in their choices of optional parameter settings (Appendix D.4). Tables 2 and 3 compare their discovery performance with the baselines. We also conduct large-scale first-principles validation on MP-20, evaluating 2,000 candidates from each model's unique-and-novel pool with VASP at a stability threshold of 0.1 eV/atom (Appendix G.4). By the standards of crystal-generation research, this is a particularly large-scale first-principles evaluation.

Table 2: MP-20 material discovery rates (%; higher is better). **Bold** highlights leading results; underlining marks runner-up results. The same convention applies to subsequent result tables.

Method	SUN@0	SSUN@0	SUN@0.1	SSUN@0.1	SUN-VASP	SSUN-VASP
DiffCSP++	7.71	7.64	37.95	37.31	24.47	24.47
SymmCD	7.34	7.18	33.11	32.02	21.56	21.52
WyFormer	6.76	6.69	33.28	32.11	19.93	19.75
SGEquiDiff	8.45	8.31	38.00	36.38	26.74	26.62
DynaCrys (ours)	**9.39**	**9.31**	**41.49**	**40.12**	**29.56**	**28.98**
DynaCrys-II (ours)	**9.18**	**9.10**	**41.32**	**39.87**	**30.63**	**30.28**

DynaCrys achieves the best SUN and SSUN rates at both thresholds on both datasets (Tables 2 and 3). Its lead over the baselines becomes more pronounced on the more challenging MPTS-52 dataset. VASP (Table 2) and MACE-MP-0 evaluation (Appendix H) further separately confirm this lead. These results demonstrate DynaCrys's leading and robust discovery performance both overall and under the additional requirement of nontrivial symmetry after relaxation.

Table 3: MPTS-52 material discovery rates (%; higher is better).

Method	SUN@0	SSUN@0	SUN@0.1	SSUN@0.1
SymmCD	5.07	4.88	21.65	20.64
WyFormer	2.05	1.81	13.58	10.57
SGEquiDiff	5.17	4.94	26.29	24.33
DynaCrys (ours)	**7.30**	**7.16**	**34.20**	**33.05**
DynaCrys-II (ours)	**6.65**	**6.58**	**31.85**	**30.85**

5.3 Validity and distributional fidelity

We also assess structural and compositional validity, distributional distances for density, the number of distinct elements, and Wyckoff degrees of freedom, and coverage recall and precision on relaxed structures (Chang et al., 2025). These metrics measure structural and chemical plausibility and agreement with the reference data. The median energy above the hull (E_{hull}) characterizes energetic quality (Tables 4 and 5). We compute distributional distances on valid structures, using the test split for MP-20, which is randomly partitioned. For MPTS-52, whose chronological split produces a substantial train–test distribution shift (Appendix G.3), we use the training split as the reference to assess how well generated structures capture the statistical properties of the training materials.

DynaCrys leads in structural validity and coverage precision on both datasets, while also achieving the smallest element-count and Wyckoff distribution distances and the lowest median energies above

Table 4: Results for validity, distributional discrepancies, and energetic properties on MP-20.

Method	Validity (%) ↑		Distribution ↓			Coverage (%) ↑		P1 (%) ↓	Median E_{hull} ↓
	Struct.	Comp.	d_ρ	$d_{N_{\text{el}}}$	d_{Wyckoff}	R	P		(eV/atom)
DiffCSP++	99.91	85.12	—	—	—	—	—	**5.10**	0.107
SymmCD	98.54	**86.27**	0.2363	0.3902	0.0722	99.78	99.24	7.60	0.115
WyFormer	99.84	82.65	0.2081	0.0753	0.0590	**99.79**	99.00	6.43	0.130
SGEquiDiff	99.50	86.13	0.1761	0.1195	0.0781	99.76	99.14	7.95	0.082
DynaCrys (ours)	**99.97**	83.36	**0.1025**	**0.0518**	**0.0340**	99.71	**99.31**	5.54	**0.072**
DynaCrys-II (ours)	**99.95**	83.30	**0.1225**	**0.0627**	**0.0344**	99.70	**99.42**	5.42	**0.060**

Table 5: Results for validity, distributional discrepancies, and energetic properties on MPTS-52.

Method	Validity (%) ↑		Distribution ↓			Coverage (%) ↑		P1 (%) ↓	Median E_{hull} ↓
	Struct.	Comp.	d_ρ	$d_{N_{\text{el}}}$	d_{Wyckoff}	R	P		(eV/atom)
SymmCD	89.25	73.53	**0.1372**	0.5610	0.1054	99.78	89.78	9.37	0.204
WyFormer	90.62	70.59	0.6904	0.2382	0.2201	**99.96**	90.56	23.80	0.323
SGEquiDiff	98.36	**79.91**	0.2400	0.1479	0.1266	99.94	96.24	10.09	0.161
DynaCrys (ours)	**99.91**	75.02	0.2618	**0.0501**	**0.0626**	99.83	**98.54**	**4.59**	**0.093**
DynaCrys-II (ours)	**99.90**	76.90	0.1779	**0.0202**	**0.0501**	99.84	**98.35**	**4.66**	**0.101**

the hull (Tables 4 and 5). On MP-20, it also leads the density-distribution comparison. These results show that DynaCrys combines the highest SUN and SSUN rates with leading structural validity, favorable energies above the hull, and close agreement with key statistics of the reference materials.

5.4 Geometric quality and computational efficiency

We assess geometric quality by comparing generated structures with their relaxed endpoints. We report the mean root-mean-square displacement (RMSD) over matched pairs and the fraction of evaluated pairs with a structural correspondence (Matched), using MatterGen's matching protocol. DynaCrys achieves the lowest RMSD and highest matching rates on both datasets (Table 6). These results show that DynaCrys generates higher-quality geometries that require smaller adjustments during relaxation and more often retain structural correspondence with their relaxed endpoints.

We also evaluate generation efficiency and the effect of fewer geometry sampling steps on MP-20. DynaCrys achieves the highest SUN@0 among the compared methods, while DynaCrys-small (trained within the same framework with a reduced hidden dimension; Appendix F) achieves a similar SUN rate and the highest generation rate for complete structures (Figure 4; timing settings in Appendix G.5). Besides, we find that halving DynaCrys's geometry sampling steps from 1,000 to 500 preserves approximately 93.8% of its SUN@0.1 rate (full results in Appendix G.6).

5.5 Ablation studies

To assess the contributions of individual design choices, we conduct the following ablation studies on MPTS-52. Table 7 reports their discovery and stability results. All ablations show worse SUN rates. A detailed comparison including additional metrics is provided in Appendix I.1 (Table 15).

Dynamic space-group generation. We train a variant without dynamic space-group diffusion: it samples the space group once and keeps it fixed throughout generation. Fixing the space group reduces SUN@0.1 by 4.77 percentage points, lowers SSUN and stability, and increases median E_{hull}. These results demonstrate the benefit of jointly revising the space group, occupations, and elements for discovering stable new crystals.

Table 6: Geometric changes under relaxation.

	MP-20		MPTS-52	
Method	RMSD ↓ (Å)	Matched ↑ (%)	RMSD ↓ (Å)	Matched ↑ (%)
DiffCSP++	0.206	95.29	—	—
SymmCD	0.377	85.27	0.523	75.64
WyFormer	0.191	95.54	0.641	45.60
SGEquiDiff	0.190	91.38	0.294	88.05
DynaCrys (ours)	**0.172**	**96.18**	**0.255**	**97.11**
DynaCrys-II (ours)	**0.164**	**96.65**	**0.239**	**96.95**

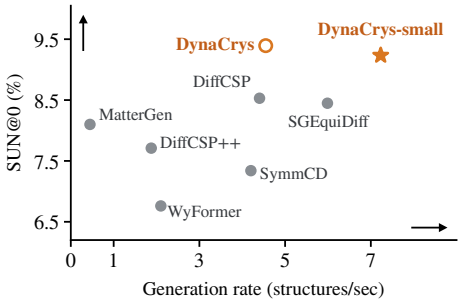


Figure 4: Generation rates of crystals on MP-20.

Group–subgroup diffusion kernel. We retain dynamic space-group generation and replace the group–subgroup kernel with an all-to-all kernel defined on a fully connected graph. The all-to-all variant better matches the reference Wyckoff distribution, but matches the element-count and density distributions less closely and yields lower SUN rates.

Table 7: Ablation studies on MPTS-52, averaged over five generation seeds. The indented Δ row is relative to w/o FiLM. Rates are in % (changes in percentage points) and energies in eV/atom.

	Discovery (%) ↑				Stability	
Method	SUN@0	SSUN@0	SUN@0.1	SSUN@0.1	Stable@0.1 ↑	Median E_{hull} ↓
DynaCrys (ours)	**7.30**	**7.16**	**34.20**	**33.05**	**51.87**	**0.093**
Ablation: w/o Dynamic Space Group	6.32	6.21	29.43	28.32	46.00	0.118
Ablation: w/o Group-subgroup Graph	6.71	6.64	31.61	30.90	50.68	0.097
Ablation: w/o FiLM	6.51	6.41	32.91	31.73	49.29	0.102
↪ **Ablation: w/o stochastic decoding (Δ)**	(-1.16)	(-1.13)	(-5.26)	(-4.67)	(-2.56)	(+0.012)

Element-head fine-tuning. We omit the element-head fine-tuning stage during training to obtain the w/o FiLM variant. The results show that removing this stage improves agreement with the reference element-count distribution, but worsens agreement with the density and Wyckoff distributions and lowers stability.

Stochastic legal decoding. Using the same checkpoint before element-head fine-tuning and the same geometry model, we compare stochastic constrained sampling with maximum a posteriori (MAP) decoding, which deterministically searches for high-probability legal configurations. MAP achieves higher validity and lower RMSD, but shows poorer agreement with the reference element-count distribution and lower stability and SUN rates.

6 Conclusion

We introduced DynaCrys, a framework for complete crystal generation in which the space group jointly evolves and is dynamically revised with Wyckoff occupations and elements. A group–subgroup diffusion kernel guides space-group updates, while a legality-constrained stochastic decoder produces crystallographically compatible protostructures for subsequent symmetry-constrained geometry generation. DynaCrys achieves state-of-the-art performance in symmetry-aware crystal generation on standard benchmarks, while supporting fast sampling. Independent MLIP evaluations and large-scale first-principles validation confirm this advantage, and ablations show the contributions of the key design choices. DynaCrys thus offers a new, strong, and more flexible framework for crystal generation. A limitation of our work is that the experimental synthesizability of the generated crystals remains unverified, as our stability assessments rely on computational screening proxies, including hull distances computed against a finite snapshot of reference structures.

AI USE STATEMENT

We used generative AI tools as assistive tools for manuscript editing and translation, literature review, research code development, mathematical checks, figure preparation, and the analysis of results. The authors reviewed the AI-assisted outputs, made the final research and writing decisions, and take full responsibility for the final manuscript and associated research artifacts.

REPRODUCIBILITY STATEMENT

The data used in this work are publicly available from their original sources. Model definitions, mathematical derivations, training settings, and sampling procedures are provided in Section 4 and Appendices A–F. Appendix G describes the datasets and evaluation protocols. Appendices H and I report independent evaluation with MACE-MP-0 and variability across generation seeds, respectively. We will also open-source the code. The required internal approval process is already ongoing.

REFERENCES

M. I. Aroyo (ed.). *International Tables for Crystallography: Space-group symmetry*, volume A. International Union of Crystallography, Chester, England, 2nd edition, December 2016. ISBN 9780470974230. doi: 10.1107/97809553602060000114. URL https://it.iucr.org/Ac/.

Jacob Austin, Daniel D. Johnson, Jonathan Ho, Daniel Tarlow, and Rianne van den Berg. Structured denoising diffusion models in discrete state-spaces. In *Advances in Neural Information Processing Systems*, volume 34, 2021.

Hartmut Bärnighausen. Group-subgroup relations between space groups: a useful tool in crystal chemistry. *MATCH Commun. Math. Comput. Chem.*, 9:139–175, 1980.

Ilyes Batatia, Philipp Benner, Yuan Chiang, Alin M. Elena, Dávid P. Kovács, Janosh Riebesell, Xavier R. Advincula, Mark Asta, Matthew Avaylon, William J. Baldwin, Fabian Berger, Noam Bernstein, Arghya Bhowmik, Filippo Bigi, Samuel M. Blau, Vlad Cărare, Michele Ceriotti, Sanggyu Chong, James P. Darby, Sandip De, Flaviano Della Pia, Volker L. Deringer, Rokas Elijošius, Zakariya El-Machachi, Edvin Fako, Fabio Falcioni, Andrea C. Ferrari, John L. A. Gardner, Mikołaj J. Gawkowski, Annalena Genreith-Schriever, Janine George, Rhys E. A. Goodall, Jonas Grandel, Clare P. Grey, Petr Grigorev, Shuang Han, Will Handley, Hendrik H. Heenen, Kersti Hermansson, Cheuk Hin Ho, Stephan Hofmann, Christian Holm, Jad Jaafar, Konstantin S. Jakob, Hyunwook Jung, Venkat Kapil, Aaron D. Kaplan, Nima Karimitari, James R. Kermode, Panagiotis Kourtis, Namu Kroupa, Jolla Kullgren, Matthew C. Kuner, Domantas Kuryla, Guoda Liepuoniute, Chen Lin, Johannes T. Margraf, Ioan-Bogdan Magdău, Angelos Michaelides, J. Harry Moore, Aakash A. Naik, Samuel P. Niblett, Sam Walton Norwood, Niamh O'Neill, Christoph Ortner, Kristin A. Persson, Karsten Reuter, Andrew S. Rosen, Louise A. M. Rosset, Lars L. Schaaf, Christoph Schran, Benjamin X. Shi, Eric Sivonxay, Tamás K. Stenczel, Christopher Sutton, Viktor Svahn, Thomas D. Swinburne, Jules Tilly, Cas van der Oord, Santiago Vargas, Eszter Varga-Umbrich, Tejs Vegge, Martin Vondrák, Yangshuai Wang, William C. Witt, Thomas Wolf, Fabian Zills, and Gábor Csányi. A foundation model for atomistic materials chemistry. *The Journal of Chemical Physics*, 163(18):184110, 11 2025. ISSN 0021-9606. doi: 10.1063/5.0297006. URL https://doi.org/10.1063/5.0297006.

Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *Computer Vision – ECCV 2020*, volume 12346 of *Lecture Notes in Computer Science*, pp. 213–229. Springer, 2020. doi: 10.1007/978-3-030-58452-8_13.

Rees Chang, Angela Pak, Alex Guerra, Ni Zhan, Nick Richardson, Elif Ertekin, and Ryan P. Adams. Space group equivariant crystal diffusion. In *Advances in Neural Information Processing Systems*, volume 38, 2025.

Bowen Deng, Peichen Zhong, KyuJung Jun, Janosh Riebesell, Kevin Han, Christopher J. Bartel, and Gerbrand Ceder. CHGNet as a pretrained universal neural network potential for charge-informed atomistic modelling. *Nature Machine Intelligence*, 5:1031–1041, 2023. doi: 10.1038/s42256-023-00716-3.

Filip Ekström Kelvinius, Oskar B. Andersson, Abhijith S. Parackal, Dong Qian, Rickard Armiento, and Fredrik Lindsten. WyckoffDiff – a generative diffusion model for crystal symmetry. In *Proceedings of the 42nd International Conference on Machine Learning (ICML)*, 2025.

Scott Fredericks, Kevin Parrish, Dean Sayre, and Qiang Zhu. PyXtal: A Python library for crystal structure generation and symmetry analysis. *Computer Physics Communications*, 261:107810, 2021. doi: 10.1016/j.cpc.2020.107810.

Rhys E. A. Goodall, Abhijith S. Parackal, Felix A. Faber, Rickard Armiento, and Alpha A. Lee. Rapid discovery of stable materials by coordinate-free coarse graining. *Science Advances*, 8(30): eabn4117, 2022. doi: 10.1126/sciadv.abn4117.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, volume 33, 2020.

Anubhav Jain, Shyue Ping Ong, Geoffroy Hautier, Wei Chen, William Davidson Richards, Stephen Dacek, Shreyas Cholia, Dan Gunter, David Skinner, Gerbrand Ceder, and Kristin A. Persson. Commentary: The Materials Project: A materials genome approach to accelerating materials innovation. *APL Materials*, 1(1):011002, 2013. doi: 10.1063/1.4812323.

Rui Jiao, Wenbing Huang, Peijia Lin, Jiaqi Han, Pin Chen, Yutong Lu, and Yang Liu. Crystal structure prediction by joint equivariant diffusion. In *Advances in Neural Information Processing Systems*, volume 36, 2023.

Rui Jiao, Wenbing Huang, Yu Liu, Deli Zhao, and Yang Liu. Space group constrained crystal generation. In *International Conference on Learning Representations (ICLR)*, 2024.

Nikita Kazeev, Wei Nong, Ignat Romanov, Ruiming Zhu, Andrey E Ustyuzhanin, Shuya Yamazaki, and Kedar Hippalgaonkar. Wyckoff transformer: Generation of symmetric crystals. In *Forty-second International Conference on Machine Learning*, 2025. URL https://openreview.net/forum?id=eFHfRQRjJo.

Daniel Levy, Siba Smarak Panigrahi, Sékou-Oumar Kaba, Qiang Zhu, Kin Long Kelvin Lee, Mikhail Galkin, Santiago Miret, and Siamak Ravanbakhsh. SymmCD: Symmetry-preserving crystal generation with diffusion models. In *International Conference on Learning Representations (ICLR)*, 2025.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=Bkg6RiCqY7.

Benjamin Kurt Miller, Ricky T. Q. Chen, Anuroop Sriram, and Brandon M. Wood. FlowMM: Generating materials with Riemannian flow matching. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*, 2024.

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

Ethan Perez, Florian Strub, Harm de Vries, Vincent Dumoulin, and Aaron Courville. FiLM: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2018.

Atsushi Togo, Kohei Shinohara, and Isao Tanaka. Spglib: a software library for crystal symmetry search. *Science and Technology of Advanced Materials: Methods*, 4(1):2384822, 2024. doi: 10.1080/27660400.2024.2384822.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30, pp. 5998–6008, 2017.

Tian Xie, Xiang Fu, Octavian-Eugen Ganea, Regina Barzilay, and Tommi Jaakkola. Crystal diffusion variational autoencoder for periodic material generation. In *International Conference on Learning Representations (ICLR)*, 2022.

Claudio Zeni, Robert Pinsler, Daniel Zügner, Andrew Fowler, Matthew Horton, Xiang Fu, Zilong Wang, Aliaksandra Shysheya, Jonathan Crabbé, Shoko Ueda, Roberto Sordillo, Lixin Sun, Jake Smith, Bichlien Nguyen, Hannes Schulz, Sarah Lewis, Chin-Wei Huang, Ziheng Lu, Yichi Zhou, Han Yang, Hongxia Hao, Jielan Li, Chunlei Yang, Wenjie Li, Ryota Tomioka, and Tian Xie. A generative model for inorganic materials design. *Nature*, 639:624–632, 2025. doi: 10.1038/s41586-025-08628-5.

Biao Zhang and Rico Sennrich. Root mean square layer normalization. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

A SYMBOLIC STATE SPACE AND LEGALITY

This appendix defines the finite state space of the symbolic model and the legal clean states used by the decoder.

A.1 SYMBOLIC STATE SPACE

Let

$$\mathcal{G} = \{G : n_G^{\text{train}} > 0\}$$

be the dataset-specific vocabulary of training-supported space groups: 168 for MP-20 and 202 for MPTS-52. Using PyXtal's space group and Wyckoff position tables (Fredericks et al., 2021), we construct the crystallographic row catalogue and associated symmetry operations for all 230 space groups, with $\mathcal{R}_G$ denoting the valid Wyckoff rows of group G. Across this catalogue there are 1,731 valid (G, r) pairs, and the row table for each space group is padded to width 27.

The symbolic state uses $M = 20$ orbit slots for MP-20 and $M = 52$ for MPTS-52. A clean state has

$$G \in \mathcal{G}, \qquad K \in \{1, \dots, M\}, \qquad \sum_{i=1}^{M} m_i = K.$$

Each active slot has $r_i \in \mathcal{R}_G$ and carries a categorical element label $Z_i \in \mathcal{Z}$, while inactive slots use separate PAD symbols. Noisy states replace the row index by the continuous row vector y_i defined in Appendix B.

A.2 CRYSTALLOGRAPHIC TABLES

The symbolic and geometry stages share the following crystallographic metadata for each valid (G, r):

Symbol	Meaning
$\mu_{G,r}$	multiplicity in the conventional cell
$\text{dof}_{G,r}$	number of free anchor coordinates
$s_{G,r}$	site-symmetry label
$e_{G,r}$	enumeration within $(G, s_{G,r})$

The valid-row mask restricts every space-group-conditional row distribution to $\mathcal{R}_G$. Multiplicity enters the atom-count constraint, while the remaining metadata parameterize the shared symmetry codebook and the geometry model. The configuration-specific capacity $u_{G,r}$ is defined below.

A.3 CAPACITY AND LEGAL CLEAN STATES

For MP-20 configuration DynaCrys, let $q_{0.99}^{\text{train}}(G, r)$ be the 99th percentile of the within-structure repetition count of row (G, r), computed on the training partition. Let $B_{G,r} = \lfloor N_{\max}/\mu_{G,r} \rfloor$ be the largest occupation compatible with the conventional-cell atom cap $N_{\max}$, set to 80 for MP-20 and 208 for MPTS-52. The resulting empirical capacity is

$$u_{G,r} = \begin{cases} 0, & r \notin \mathcal{R}_G, \\ 1, & \text{dof}_{G,r} = 0 \text{ or } q_{0.99}^{\text{train}}(G, r) \text{ is unavailable}, \\ \max\{1, \min(\lceil q_{0.99}^{\text{train}}(G, r) \rceil, B_{G,r})\}, & \text{otherwise}. \end{cases} \tag{10}$$

MP-20 configuration DynaCrys-II and both MPTS-52 configurations use the multiplicity-based capacity rule in Appendix D.4, which also records the configuration differences.

Write a padded clean state as $x_0 = (G, K, A, R, Z)$, where $A \subseteq \{1, \dots, M\}$ is the active-slot set and $R = (r_i)_{i \in A}$, $Z = (Z_i)_{i \in A}$. Define the number of active occupations assigned to row r by

$$n_r(A, R) = \sum_{i \in A} \mathbf{1}[r_i = r].$$

For a chosen capacity table u, the legal clean-state set in Eq. (6) is

$$\mathcal{X}_{\text{legal}} \equiv \mathcal{X}_{\text{legal}}(u) = \left\{ \begin{array}{ll} (G, K, A, R, Z): & G \in \mathcal{G}, \quad |A| = K, \\ (r_i, Z_i) \in \mathcal{R}_G \times \mathcal{Z} & \forall i \in A, \\ n_r(A, R) \le u_{G,r} & \forall r \in \mathcal{R}_G, \\ \sum_{i \in A} \mu_{G,r_i} \le N_{\max} & \end{array} \right\}. \tag{11}$$

B Forward Corruption Kernels

The symbolic forward process is a family of marginal kernels $q_t(x_t \mid x_0)$. At each noise level it samples

$$G_t \longrightarrow K_t \longrightarrow \text{survivor–birth–PAD partition} \longrightarrow \{(m_{i,t}, Z_{i,t}, y_{i,t})\}_{i=1}^{M}.$$

The same family is used for training-time corruption and reverse-time re-noising.

B.1 Training priors

All reset and birth distributions are estimated on the training partition. Writing n_G, $n_{G,r}$, and n_Z for the corresponding counts,

$$\pi_G(G) = \frac{n_G}{\sum_{G' \in \mathcal{G}} n_{G'}}, \qquad \pi_Z(Z) = \frac{n_Z}{\sum_{Z' \in \mathcal{Z}} n_{Z'}}. \tag{12}$$

The row prior is smoothed within the valid vocabulary of each group:

$$\pi_R(r \mid G) = \frac{n_{G,r} + 1}{\sum_{r' \in \mathcal{R}_G} (n_{G,r'} + 1)}, \qquad r \in \mathcal{R}_G. \tag{13}$$

For the orbit count, let $\widehat{p}_{\text{train}}(K \mid G)$ and $\widehat{p}_{\text{train}}(K)$ be the space-group-conditional and global empirical distributions. We use

$$\pi_K(K \mid G) = (1 - \rho_G)\widehat{p}_{\text{train}}(K \mid G) + \rho_G \widehat{p}_{\text{train}}(K), \qquad \rho_G = \frac{\alpha_K}{n_G + \alpha_K}, \quad \alpha_K = 20. \tag{14}$$

Here $\alpha_K = 20$ is a smoothing pseudo-count controlling shrinkage toward the global orbit-count distribution.

B.2 Space-group graph and kernel

The graph on $\mathcal{G}$ contains the maximal translationengleiche and klassengleiche group–subgroup relations, with edge directions removed and self-relations omitted. Let A be its adjacency matrix, $\widetilde{A} = A + I$, $D = \operatorname{diag}(\widetilde{A}\mathbf{1})$, and

$$W_{\mathcal{G}} = D^{-1}\widetilde{A}, \qquad Q = W_{\mathcal{G}} - I, \qquad R_s = \exp(\beta_{\max} s Q), \quad \beta_{\max} = 4, \tag{15}$$

where $s = t/T$. The space-group channel is

$$q_t^G(G_t \mid G_0) = \lambda_G(s) R_s(G_t \mid G_0) + \big[1 - \lambda_G(s)\big]\pi_G(G_t). \tag{16}$$

B.3 Orbit-count and slot channels

Given the sampled group, the orbit count follows

$$q_t^K(K_t \mid K_0, G_t) = \lambda_K(s)\delta_{K_0}(K_t) + \big[1 - \lambda_K(s)\big]\pi_K(K_t \mid G_t). \tag{17}$$

The number of surviving clean occupations is

$$S_t \sim \text{Binomial}(\min(K_0, K_t),\, \alpha_S(s)). \tag{18}$$

The S_t survivors are matched to distinct clean occupations sampled uniformly without replacement. Together with $K_t - S_t$ births, they are placed uniformly among the M slots; the remaining slots are PAD.

Let $c_{G,r}$ and c_{pad} be the row and padding vectors of the shared codebook in Appendix C. For each slot,

$$y_{i,t} = \mu_{i,t} + \sigma_y(s)\epsilon_i, \qquad \epsilon_i \sim \mathcal{N}(0, I_d), \tag{19}$$

with

$$\mu_{i,t} = \begin{cases} \sqrt{\bar{\alpha}_y(s)}\, c_{G_0,r_i^0}, & \text{survivor matched to } r_i^0, \\ c_{G_t,r_i^{\mathrm{birth}}}, \quad r_i^{\mathrm{birth}} \sim \pi_R(\cdot \mid G_t), & \text{birth}, \\ c_{\mathrm{pad}}, & \text{PAD}. \end{cases} \tag{20}$$

At survivor slots, elements follow

$$q_t^Z(Z_{i,t} \mid Z_i^0) = \alpha_Z(s)\delta_{Z_i^0}(Z_{i,t}) + \big[1 - \alpha_Z(s)\big]\pi_Z(Z_{i,t}). \tag{21}$$

Birth elements are drawn from π_Z, while PAD slots carry the element PAD symbol.

B.4 Schedules and terminal prior

The discrete channels share a quadratic retention schedule, and the row channel uses a cosine signal schedule:

$$\begin{aligned} \lambda_G(s) = \lambda_K(s) = \alpha_S(s) = \alpha_Z(s) = (1-s)^2, \\ \bar{\alpha}_y(s) = \cos^2\Big(\frac{\pi s}{2}\Big), \qquad \sigma_y(s) = \sqrt{1 - \bar{\alpha}_y(s)}. \end{aligned} \tag{22}$$

Here $\bar{\alpha}_y(s)$ is the retained row-signal power and $\sigma_y(s)$ is the row-noise scale. The training horizon is $T = 1000$; at generation time, the reverse sampler operates at 200 noise levels arranged in descending order along this schedule.

At $s = 1$, all clean-state retention terms vanish. The terminal state can therefore be sampled directly as

$$\begin{aligned} &G_T \sim \pi_G, \qquad K_T \sim \pi_K(\cdot \mid G_T), \qquad A_T \sim \mathrm{Unif}\{A \subseteq \{1,\dots,M\} : |A| = K_T\}, \\ &r_{i,T} \sim \pi_R(\cdot \mid G_T), \qquad Z_{i,T} \sim \pi_Z, \qquad y_{i,T} = c_{G_T,r_{i,T}} + \epsilon_i \quad (i \in A_T). \end{aligned} \tag{23}$$

Here $m_{i,T} = \mathbf{1}[i \in A_T]$. On PAD slots, $y_{i,T} = c_{\mathrm{pad}} + \epsilon_i$, and $Z_{i,T}$ is the element PAD symbol. This distribution contains no dependence on x_0 and is the starting distribution of the reverse sampler.

C Shared Symmetry Codebook

The codebook maps space groups and space-group-specific Wyckoff rows into a common vector space. It is constructed once, frozen, and shared by the symbolic and geometry stages.

C.1 Space-group and row representations

For a symmetry operation $g = (R \mid \tau)$ in the fractional-coordinate basis, define

$$\phi(g) = [\mathrm{vec}(R),\ \tau \bmod 1,\ \det R,\ \mathrm{tr}\, R] \in \mathbb{R}^{14}. \tag{24}$$

A shared operation encoder e_ϕ maps these features to $\mathbb{R}^d$. Let S_G denote the operation set of space group G. A learned map f_G produces the space-group embedding from the pooled operation features:

$$c_G = f_G[\mathrm{mean}_{g \in S_G}\, e_\phi(g) \parallel \max_{g \in S_G} e_\phi(g)]\,. \tag{25}$$

For row r, let $\mathcal{C}_{G,r}$ be its orbit-generating coset representatives and $\mathcal{H}_{G,r}$ the stabilizer of its representative. A learned map f_R produces the row embedding from the space-group embedding, operation-derived features, and row metadata:

$$\begin{aligned} c_{G,r} = f_R\Big[&c_G,\ \mathrm{mean}_{g \in \mathcal{C}_{G,r}}\, e_\phi(g),\ \mathrm{mean}_{g \in \mathcal{H}_{G,r}}\, e_\phi(g), \\ &\mathrm{emb}(\mu_{G,r}),\ \mathrm{emb}(\mathrm{dof}_{G,r}),\ \mathrm{emb}(s_{G,r}),\ \mathrm{emb}(e_{G,r})\Big]. \end{aligned} \tag{26}$$

Here $e_{G,r}$ enumerates rows within the same $(G, \text{site-symmetry})$ class. The codebook width is $d = 128$, and each metadata embedding has width 16. A separately initialized c_{pad} represents padded rows.

C.2 PRETRAINING AND FREEZING

Pretraining has two stages. The first fits deterministic crystallographic labels from operation-derived representations. From c_G, the model predicts the space-group type, crystal system, centering, and point group; an operation-only row representation predicts the number of free coordinates and the site-symmetry label. The sum of these cross-entropies is denoted by $\mathcal{L}_{\mathrm{struct}}$.

The second stage freezes the symmetry-operation encoder and the space-group encoder and refines the modules that construct the row embeddings. A row center is corrupted with the row-channel cosine schedule in Eq. (22),

$$\widetilde{c}_{G,r} = \sqrt{\bar{\alpha}_y(s)}\, c_{G,r} + \sqrt{1-\bar{\alpha}_y(s)}\, \epsilon, \qquad \epsilon \sim \mathcal{N}(0, I_d),$$

and retrieved among the valid rows of the same group with a cosine-softmax loss

$$\mathcal{L}_{\mathrm{retr}} = -\log \frac{\exp(\cos(\widetilde{c}_{G,r}, c_{G,r})/\tau_{\mathrm{retr}})}{\sum_{r' \in \mathcal{R}_G} \exp(\cos(\widetilde{c}_{G,r}, c_{G,r'})/\tau_{\mathrm{retr}})},$$

where τ_{retr} is the retrieval temperature. Distinct rows within a group are separated by

$$\mathcal{L}_{\mathrm{sep}} = \operatorname*{mean}_{\substack{G,\ r \neq r' \\ r,r' \in \mathcal{R}_G}} \left[\mathrm{ReLU}(\cos(c_{G,r}, c_{G,r'}) - m)^2\right]. \tag{27}$$

The second-stage objective is

$$\mathcal{L}_{\mathrm{cb}} = \mathcal{L}_{\mathrm{retr}} + \lambda_{\mathrm{sep}} \mathcal{L}_{\mathrm{sep}}. \tag{28}$$

The shared training settings are summarized below.

Setting	**Value**
operation / codebook width	128 / 128
metadata width	16
steps per stage	3,000
optimizer	AdamW (Loshchilov & Hutter, 2019) with cosine decay
learning rate	3×10^{-3}
stage-1 / stage-2 weight decay	0 / 0.01
retrieval temperature / noise draws	0.1 / 4
separation margin m / weight λ_{sep}	0.5 / 1

After training, every nonzero vector is normalized as

$$c \leftarrow \sqrt{d}\, \frac{c}{\|c\|_2}, \tag{29}$$

so each vector has unit root-mean-square coordinate magnitude. The resulting c_G, $c_{G,r}$, and c_{pad} tables are frozen.

C.3 USE IN DYNACRYS

The frozen c_{G_t} represents the current noisy space group in the symbolic denoiser. The row vectors $c_{G,r}$ serve both as centers of the continuous row channel and as the tied vectors of the space-group-conditional row head. The accepted c_G and $c_{G,r}$ also condition the geometry model. All reported DynaCrys configurations use the same frozen codebook.

D SYMBOLIC DENOISER AND TRAINING

This appendix specifies the set denoiser, its prediction heads, and the objective summarized in Section 4.2.

D.1 PERMUTATION-EQUIVARIANT SET DENOISER

The noisy state is encoded as one global token and M orbit-slot tokens:

$$h^{\text{in}}_{\text{global}} = \text{MLP}[c_{G_t}, e_{K_t}, e_t]\,, \qquad h^{\text{in}}_i = \text{MLP}[e_{m_i}, e_{Z_i}, y_i, e_t, c_{G_t}]\,. \tag{30}$$

Here e_t is a sinusoidal time embedding passed through an MLP. A Transformer encoder processes the tokens without slot positional encodings. Its slot outputs are therefore equivariant to permutations of the orbit occupations, while the global output is invariant to those permutations.

Linear heads on the global output predict G_0 and K_0. A slot-wise active head predicts whether each output corresponds to an occupied orbit or padding. The remaining slot heads predict its Wyckoff row and element.

D.2 ROW AND ELEMENT HEADS

For a candidate space group G, the row head classifies only over $\mathcal{R}_G$. Its logits are tied to the shared codebook with RMS normalization (Zhang & Sennrich, 2019):

$$\begin{aligned} \ell^{(G)}_{i,r} &= \frac{\widehat{q}_i^{\top}\widehat{c}_{G,r}}{\tau} + b_{G,r}, \\ \widehat{q}_i &= \text{normalize}(\text{RMSNorm}(W_q h_i) + g_y(t)\lambda_y\,\text{RMSNorm}(y_i))\,, \end{aligned} \tag{31}$$

where $\widehat{c}_{G,r} = \text{normalize}(c_{G,r})$, τ and λ_y are learned scalars, $b_{G,r}$ is a row bias, $s = t/T$, and $g_y(t) = \sqrt{\bar{\alpha}_y(s)}$. The invalid-row logits are masked before the row-wise softmax. During training, the clean space group G_0 determines the row vocabulary $\mathcal{R}_{G_0}$; during sampling, the row head is evaluated over $\mathcal{R}_G$ for each proposed clean space group G.

The element head conditions its logits on the selected (G, r) through feature-wise linear modulation (FiLM) (Perez et al., 2018):

$$\ell^{Z}_i(G, r) = [1 + \gamma_{\text{raw}}(c_{G,r})] \odot W_Z h_i + \beta(c_{G,r}), \tag{32}$$

where a shared MLP produces γ_{raw} and β from the row code. The modulation is initialized at the identity. Its regularizer is

$$\mathcal{L}_{\text{FiLM}} = \lambda_{\text{FiLM}}\,\frac{1}{|\mathcal{Z}|}\,\underset{(G,r)\ \text{valid}}{\text{mean}}\left(\|\gamma_{\text{raw}}(c_{G,r})\|_2^2 + \|\beta(c_{G,r})\|_2^2\right). \tag{33}$$

Here $|\mathcal{Z}|$ is the element vocabulary size.

D.3 SET OBJECTIVE AND ASSIGNMENT

The clean-state prediction loss is

$$\begin{aligned} \mathcal{L}_{\text{sym}} = \mathcal{L}_G + w_K\mathcal{L}_K + w_m\mathcal{L}_{\text{mask}} + w_r\mathcal{L}_r + w_Z\mathcal{L}_Z \\ + s_{\text{aux}}\left(w_{\text{cap}}\mathcal{L}_{\text{cap}} + w_N\mathcal{L}_N + w_{\text{comp}}\mathcal{L}_{\text{comp}}\right). \end{aligned} \tag{34}$$

The global losses are the clean-target cross-entropies $\mathcal{L}_G = -\log p_\theta(G_0 \mid x_t)$ and $\mathcal{L}_K = -\log p_\theta(K_0 \mid x_t)$ for the predicted space-group and orbit-count distributions, respectively. For the slot losses, a Hungarian assignment aligns the M predictions with the K_0 clean occupations and $M - K_0$ padding targets. Writing $\pi(j)$ for the predicted slot matched to target j and $\mathfrak{S}_M$ for the set of permutations of the M slots, the optimal assignment and its costs are

$$\begin{aligned} &\pi^\star = \arg\min_{\pi\in\mathfrak{S}_M}\left[\sum_{j=1}^{K_0} C_{\text{occ}}\big(\pi(j), j\big) + \sum_{j=K_0+1}^{M} C_{\text{pad}}\big(\pi(j)\big)\right], \\ &C_{\text{occ}}(i,j) = -w_m\log p_i^{\text{act}}(1 \mid x_t) - w_r\log p_i(r_j \mid G_0, x_t) - w_Z\log p_i(Z_j \mid G_0, r_j, x_t), \\ &C_{\text{pad}}(i) = -w_m\log p_i^{\text{act}}(0 \mid x_t). \end{aligned} \tag{35}$$

Here C_{occ} is the cost of assigning a predicted slot to a clean occupation, whereas C_{pad} is the cost of assigning it to a padding target. After assignment, $\mathcal{L}_{\text{mask}}$ is the binary cross-entropy between

each slot's active-state prediction and its matched target. The losses $\mathcal{L}_r$ and $\mathcal{L}_Z$ are cross-entropies between the predicted row and element distributions and their respective clean targets over the matched occupied slots.

Let $\widehat{m}_i = p_i^{\mathrm{act}}(1 \mid x_t)$ denote the predicted probability that slot i is active. The auxiliary losses are computed from the expected row occupations $\widehat{n}_r = \sum_i \widehat{m}_i p_i(r \mid G_0)$:

$$\begin{aligned}
&\mathcal{L}_{\mathrm{cap}} = \sum_r \mathrm{ReLU}(\widehat{n}_r - u_{G_0,r})^2\,, \\
&\mathcal{L}_N = (\widehat{N} - N_{\mathrm{sym},0})^2, \\
&\text{where} \quad \widehat{N} = \sum_i \widehat{m}_i \sum_r p_i(r \mid G_0)\mu_{G_0,r}.
\end{aligned} \tag{36}$$

$$\begin{aligned}
&\mathcal{L}_{\mathrm{comp}} = \sum_a (\widehat{C}_a - C_{0,a})^2, \\
&\text{where} \quad \widehat{C}_a = \sum_i \widehat{m}_i \sum_r p_i(r \mid G_0)\mu_{G_0,r} p_i(a \mid G_0, r).
\end{aligned} \tag{37}$$

Here $u_{G_0,r}$ is the permitted occupation count, $\mu_{G_0,r}$ is the Wyckoff multiplicity, and $\widehat{N}$ and $N_{\mathrm{sym},0}$ are the expected and clean conventional-cell atom counts, respectively. Likewise, $\widehat{C}_a$ and $C_{0,a}$ are the expected and clean multiplicity-weighted counts of element a. The decoder in Appendix E separately enforces the row-capacity limits and its fixed atom-count cap.

D.4 Training settings

The symbolic denoiser is trained with Eq. (34). The row-conditioned element head is then refined while all other parameters are frozen, using

$$\mathcal{L}_{\mathrm{refine}} = \mathcal{L}_{\mathrm{sym}} + \mathcal{L}_{\mathrm{FiLM}}. \tag{38}$$

The reported configurations use dataset-specific capacity rules and symbolic training schedules. Under the multiplicity-based rule, a valid row has capacity

$$u_{G,r}^{\mathrm{mult}} = \begin{cases} 1, & \mathrm{dof}_{G,r} = 0, \\ \max\{1, \lfloor N_{\max}/\mu_{G,r} \rfloor\}, & \mathrm{dof}_{G,r} > 0. \end{cases}$$

On MP-20, DynaCrys uses training-set 99th-percentile row capacities, with training budgets of 100,000 backbone steps followed by 20,000 element-head refinement steps. DynaCrys-II uses multiplicity-based capacities and budgets of 150,000 and 40,000 steps, respectively. On MPTS-52, both models use multiplicity-based capacities: DynaCrys has budgets of 150,000 backbone and 40,000 refinement steps, and DynaCrys-II has budgets of 100,000 and 20,000 steps. For each stage, we select the checkpoint with the lowest validation loss.

Table 8 lists the common architecture and optimization settings, with the dataset-specific orbit-slot counts.

Table 8: Symbolic-denoiser architecture and optimization settings.

Architecture	**Value**	**Optimization**	**Value**
hidden width	512	optimizer	AdamW
layers / attention heads	8 / 8	peak learning rate	3×10^{-4}
FFN width / dropout	$4d_h$ / 0	weight decay / grad. clip	0.01 / 1.0
time / K / mask / element emb.	128 / 32 / 16 / 64	batch size	256
orbit slots M (MP-20 / MPTS-52)	20 / 52	learning-rate schedule	2,000-step warmup; cosine to $0.05\times$ peak
FiLM hidden width	128		
core loss weights	$w_K = w_m = w_r = w_Z = 1$	auxiliary ramp	5,000 steps
auxiliary weights	$w_{\mathrm{cap}} = 0.1$, $w_N = w_{\mathrm{comp}} = 0.01$	λ_{FiLM}	10^{-4}

E LEGALITY-CONSTRAINED DECODING

E.1 CONSTRAINED TARGET

Fix a noisy state x_t, and abbreviate the denoiser outputs in Eq. (4) as $p_G(G)$, $p_K(K)$, active probabilities π_i, row distributions $p_i(r \mid G)$, and element conditionals $p_i(Z \mid G, r)$. For a clean symbolic state $x_0 = (G, K, A, \mathbf{r}, \mathbf{Z})$ with active-slot set $A \subseteq \{1, \dots, M\}$, the factorized joint distribution is

$$p_\theta^{\text{fac}}(x_0 \mid x_t) = p_G(G)\, p_K(K) \prod_{i \in A} \pi_i\, p_i(r_i \mid G)\, p_i(Z_i \mid G, r_i) \prod_{i \notin A} (1 - \pi_i). \tag{39}$$

Its joint score is $S_\theta(x_0 \mid x_t) = \log p_\theta^{\text{fac}}(x_0 \mid x_t)$. Given the legal set $\mathcal{X}_{\text{legal}}$ from Section 4.3, the stochastic decoder targets the constrained distribution

$$\widetilde{p}_\theta(x_0 \mid x_t) = \frac{p_\theta^{\text{fac}}(x_0 \mid x_t) \mathbf{1}[x_0 \in \mathcal{X}_{\text{legal}}]}{\sum_{x \in \mathcal{X}_{\text{legal}}} p_\theta^{\text{fac}}(x \mid x_t)}. \tag{40}$$

For the decoder comparison, we use a constrained maximum a posteriori (MAP) comparator that retains the top five candidates under p_G and p_K, respectively, and performs deterministic constrained search over the resulting G–K pairs.

E.2 ELEMENTARY-SYMMETRIC STOCHASTIC PROPOSAL

Write $\omega_i = \pi_i/(1 - \pi_i) = \exp(a_i)$, where a_i is the active logit. The elementary symmetric polynomial

$$e_k(\boldsymbol{\omega}) = \sum_{\substack{A \subseteq \{1, \dots, M\} \\ |A| = k}} \prod_{i \in A} \omega_i \tag{41}$$

normalizes the distribution over size-k active sets. It is evaluated by the recursion

$$F_i(k) = F_{i+1}(k) + \omega_i F_{i+1}(k-1), \qquad F_{M+1}(0) = 1, \quad F_{M+1}(k > 0) = 0, \tag{42}$$

for which $F_1(k) = e_k(\boldsymbol{\omega})$. The same table gives the conditional inclusion probability

$$\Pr(i \in A \mid k \text{ active slots remain}) = \frac{\omega_i F_{i+1}(k-1)}{F_i(k)}. \tag{43}$$

The proposal draws $G \sim p_G$, draws K with probability proportional to $p_K(K) e_K(\boldsymbol{\omega})$, draws A from

$$q(A \mid K) = \frac{\prod_{i \in A} \omega_i}{e_K(\boldsymbol{\omega})} \quad \text{for } |A| = K, \tag{44}$$

and then draws each $r_i \sim p_i(\cdot \mid G)$.

E.3 EXACTNESS

We prove Proposition 1. Let $s = (G, K, A, \mathbf{r})$ denote the symbolic skeleton, and let $\mathcal{S}_{\text{legal}}$ be the projection of $\mathcal{X}_{\text{legal}}$ onto these variables. Because $\sum_{Z \in \mathcal{Z}} p_i(Z \mid G, r_i) = 1$ for every occupied slot, marginalizing Eq. (39) over $\mathbf{Z}$ gives

$$p_\theta^{\text{skel}}(s \mid x_t) = c_0\, p_G(G)\, p_K(K) \prod_{i \in A} \omega_i\, p_i(r_i \mid G), \qquad c_0 = \prod_{i=1}^{M} (1 - \pi_i). \tag{45}$$

Let $Z_K = \sum_k p_K(k) e_k(\boldsymbol{\omega})$. One proposal attempt has density

$$\begin{aligned} q(s) &= p_G(G) \frac{p_K(K) e_K(\boldsymbol{\omega})}{Z_K} \frac{\prod_{i \in A} \omega_i}{e_K(\boldsymbol{\omega})} \prod_{i \in A} p_i(r_i \mid G) \\ &= \frac{p_\theta^{\text{skel}}(s \mid x_t)}{Z_K c_0}. \end{aligned} \tag{46}$$

The elementary-symmetric factors therefore cancel, and the proposal is proportional to the unconstrained skeleton marginal. Conditioning an attempt on $s \in \mathcal{S}_{\text{legal}}$ cancels the remaining constant:

$$\Pr(s \mid s \in \mathcal{S}_{\text{legal}}) = \frac{p_\theta^{\text{skel}}(s \mid x_t)\mathbf{1}[s \in \mathcal{S}_{\text{legal}}]}{\sum_{s' \in \mathcal{S}_{\text{legal}}} p_\theta^{\text{skel}}(s' \mid x_t)}. \tag{47}$$

After acceptance, drawing each element from $p_i(\cdot \mid G, r_i)$ restores the full joint distribution in Eq. (40). Thus the accepted output is an exact sample from the constrained target without evaluating its global normalizer.

E.4 Sampling procedure

For each reverse step, the elementary-symmetric table is computed once. For at most R attempts, the decoder draws $(G, K, A, \mathbf{r})$ from the proposal above and checks whether the resulting skeleton lies in $\mathcal{S}_{\text{legal}}$. On acceptance, it draws $\mathbf{Z}$ from the element conditionals and returns the complete clean symbolic state; otherwise it redraws the entire skeleton. If none of the R attempts is accepted, the decoder reports a rejection to the reverse sampler.

Algorithm 1 summarizes the generation path with accepted symbolic proposals.

Algorithm 1 DynaCrys generation

Require: Symbolic denoiser, forward kernels q_t, sampled levels $t_1 > \cdots > t_J$, geometry model
1: Draw x_{t_1} from the terminal prior.
2: **for** $j = 1, \ldots, J$ **do**
3: Predict the clean-state distributions from x_{t_j}.
4: Decode a clean symbolic state $\hat{x}_0^{(t_j)}$ using the procedure above.
5: **if** $j < J$ **then**
6: Draw $x_{t_{j+1}} \sim q_{t_{j+1}}(\cdot \mid \hat{x}_0^{(t_j)})$.
7: **end if**
8: **end for**
9: Remove inactive slots and slot order from $\hat{x}_0^{(t_J)}$ to obtain P.
10: Generate the lattice and orbit anchors conditioned on P (Appendix F).
11: Expand the anchors through the Wyckoff operations of P and return the crystal.

F Symmetry-Constrained Geometry Generation

The geometry model receives an accepted protostructure $P = (G, K, \{(Z_i, r_i)\}_{i=1}^K)$ from the symbolic stage and generates a lattice together with the free fractional coordinates of its occupied Wyckoff orbits. Its symmetry-constrained diffusion of the lattice representation and orbit-anchor coordinates builds on DiffCSP++ (Jiao et al., 2024). Here G determines the lattice subspace, each r_i determines the orbit chart and expansion operators, and the denoiser receives the corresponding element and shared-codebook embeddings.

F.1 Geometry denoiser

The denoiser is an atom-level message-passing network over the expanded unit cell. Periodicity enters through Fourier features of fractional-coordinate differences and the current lattice representation. The network also receives the diffusion time and protostructure conditioning. It produces a coordinate denoising field for every atom and a pooled lattice-noise prediction.

F.2 Lattice diffusion under crystal-family constraints

Let $b \in \mathbb{R}^6$ be the coordinates of $\log\sqrt{LL^\top}$ in an orthonormal basis for symmetric matrices. This rotation-invariant representation is projected onto the crystal-family subspace associated with G by the affine map $\Pi_G(b) = \Pi_G^{\text{lin}} b + d_G$. The forward process is

$$b_t = \Pi_G\big(\sqrt{\bar{\alpha}_t}\, b_0 + \sqrt{1-\bar{\alpha}_t}\, \epsilon\big), \qquad \epsilon_G = \Pi_G^{\text{lin}} \epsilon, \quad \epsilon \sim \mathcal{N}(0, I), \tag{48}$$

and the lattice head minimizes

$$\mathcal{L}_{\text{lat}} = \mathbb{E}\left[\|\epsilon_G - \widehat{\epsilon}_\theta(b_t, X_t, P, t)\|_2^2\right]. \tag{49}$$

Projection is applied after every noising and reverse update. With the usual DDPM notation (Ho et al., 2020) and $z_t \sim \mathcal{N}(0, I_6)$, the reverse step is

$$b_{t-1} = \Pi_G\left[\frac{1}{\sqrt{\alpha_t}}\left(b_t - \frac{1-\alpha_t}{\sqrt{1-\bar{\alpha}_t}}\,\widehat{\epsilon}_\theta\right) + \sigma_{t,\text{lat}}\, z_t\right], \qquad \sigma_{t,\text{lat}}^2 = \beta_t \frac{1-\bar{\alpha}_{t-1}}{1-\bar{\alpha}_t}. \tag{50}$$

F.3 Anchor-coordinate score matching

Each occupied Wyckoff orbit is parameterized by one anchor coordinate. Its Wyckoff affine operators deterministically expand that anchor into the full orbit. During training, coordinate noise is mapped to the free anchor coordinates, where the denoising target is evaluated. The predicted field is mapped to the same coordinates for the loss. During sampling, the predicted field is mapped in the same way before the updated anchors are re-expanded.

Let $u^*_{\text{wn}}(X_t, X_0; \sigma_t)$ denote the wrapped-normal denoising target in the orbit-anchor coordinates under the adopted noise normalization, and let $\Pi_{\mathcal{A}(P)}$ map each predicted atomic field to the corresponding anchor coordinates defined by P. The coordinate objective is

$$\mathcal{L}_{\text{coord}} = \mathbb{E}\left[\left\|\Pi_{\mathcal{A}(P)}\widehat{u}_\theta(X_t, b_t, P, t) - u^*_{\text{wn}}(X_t, X_0; \sigma_t)\right\|_2^2\right]. \tag{51}$$

The reverse process uses a coordinate corrector followed by a joint predictor step for the lattice and coordinates. After each update, orbit expansion reconstructs the complete unit cell, so every intermediate structure obeys the generated space group.

F.4 Architecture and training settings

A separate geometry model is trained on the training split of each dataset, using conventional-cell lattices and orbit anchors. The symbolic configurations within a dataset share its geometry-model weights. Both datasets use the common geometry architecture and training settings in Table 9; the smaller variant, DynaCrys-small, is evaluated on MP-20 in the efficiency comparison.

Table 9: Geometry model architecture and training settings.

Architecture	**Value**	**Optimization / diffusion**	**Value**
DynaCrys (ours) width / layers	512 / 6	optimizer	AdamW
DynaCrys-small (ours) width / layers	256 / 4	learning rate / weight decay	10^{-3} / 10^{-2}
		batch size / grad. clip	128 / 1.0
		warmup / training-step budget	500 / 60,000
		lattice schedule	cosine, 1,000 levels
		coordinate noise range	$[0.005, 0.5]$, 1,000 levels
		corrector coefficient	10^{-5}

G Evaluation Protocol

G.1 Datasets, relaxation, and references

We evaluate MP-20 (Xie et al., 2022) and MPTS-52 (Chang et al., 2025) separately, using each dataset's training partition for novelty and its test partition for coverage. Discovery rates in the main comparisons use a candidate budget of approximately 10,000 per evaluated set. The MP-20 geometry-step comparison uses the 10,000 generation attempts per seed as its denominator (Appendix G.6).

CHGNet (Deng et al., 2023) relaxes the generated structures. Structures from both datasets are also independently relaxed with MACE-MP-0 (Batatia et al., 2025); the two engines define separate evaluation scales. Energies above the hull are evaluated against the fixed, MP2020-corrected Materials Project (Jain et al., 2013) DFT snapshot dated 2023-02-07. The stability thresholds are $\tau = 0$ and 0.1 eV/atom.

G.2 Discovery and symmetry

Structural uniqueness is evaluated within each generated set and novelty against the corresponding training partition, using pymatgen's `StructureMatcher` (Ong et al., 2013) with its default matching tolerances. SUN@τ counts candidates that are unique, novel, and satisfy $E_{\text{hull}} \leq \tau$, divided by the candidate budget. SSUN@τ further requires a successfully detected non-$P1$ space group after relaxation. Space groups are detected with spglib (Togo et al., 2024) via `SpacegroupAnalyzer`, using `symprec`$= 0.1$ Å and an angular tolerance of $5°$. Failed symmetry identifications do not enter the SSUN numerator. Stable@τ in the ablation tables counts all stable candidates over the same candidate budget, without the uniqueness and novelty requirements.

The space-group distribution in Section 5.1 instead uses as-generated coordinates and the full MP-20 dataset as the reference, with the same symmetry tolerances. The $P1$ fractions in Tables 4 and 5 use CHGNet-relaxed structures after conventional-cell standardization: their numerators count $P1$ structures and their denominators count successful symmetry identifications. These rates are computed within each run before averaging.

G.3 Validity, distributional fidelity, and geometric quality

Tables 4 and 5 describe CHGNet-relaxed structures. Structural validity tests the minimum interatomic separation; compositional validity uses the CDVAE-compatible oxidation-state and electronegativity checks, including the alloy convention. Distributional distances use structurally and compositionally valid relaxed structures, with the MP-20 test partition and the MPTS-52 training partition as references. We use Wasserstein distances for mass density and the number of distinct elements, and Jensen–Shannon divergence in bits for the pooled distribution of Wyckoff degrees of freedom, counting each occupied orbit once. Wyckoff positions are identified at `symprec`$= 0.01$ Å and an angular tolerance of $5°$.

Train–test distribution distances are reported in Table 10. Across all three statistics, the train–test differences are negligible for MP-20 but pronounced for MPTS-52, consistent with the random split of MP-20 and the chronological split of MPTS-52.

Table 10: Distribution distances between the training and test splits, computed using all available reference features.

Dataset	d_ρ (g/cm^3)	$d_{N_{\text{el}}}$	d_{Wyckoff} (bits)
MP-20	0.048616	0.011125	0.000130
MPTS-52	0.892426	0.633220	0.026667

Coverage recall and precision use joint structural and compositional fingerprint matching against the test partition (Xie et al., 2022). On MP-20, validity and coverage retain their evaluated-set denominators. On MPTS-52, structural validity, compositional validity, and coverage precision are normalized by the 10,000-candidate budget. Coverage recall always uses the test reference. Median E_{hull} summarizes the relaxed structures with evaluable hull energies.

For RMSD, we pair each as-generated structure with its CHGNet-relaxed endpoint using the same MatterGen-derived matching procedure (Zeni et al., 2025) on both datasets. The generated structure is refined at `symprec`$= 0.01$ Å and $5°$ angular tolerance, converted to a primitive cell, and LLL-reduced; the relaxed endpoint is used as stored. Primitive-cell matching and supercell attempts are enabled, with scaling disabled and subsets disallowed. We report the mean root-mean-square displacement (RMSD, in Å) over matched structure pairs, and Matched is the percentage of all evaluated pairs for which correspondence is found.

G.4 First-principles validation on MP-20

For each method, we form a unique-and-novel pool using its CHGNet-relaxed structures and select the first 2,000 candidates in a fixed random ordering of that pool. Multi-run methods use generation seed 0; for the other baselines, we use their publicly available generated samples. VASP (version 6.4.2)

starts from the corresponding original generated structures, using the Materials Project relaxation settings supplied by the standard `MPRelaxSet` input generator. Energies receive the MP2020 compatibility corrections before evaluation against the same reference hull. The validation window is selected without ranking candidates by their CHGNet energies.

Let N_{gen} be the original candidate population, N_{UN} the size of its CHGNet unique-and-novel pool, and n_{stable} the number of converged compounds in the VASP window with $E_{\text{hull}} \leq 0.1$ eV/atom. Compounds must contain at least two distinct elements. The reported rate is

$$\text{SUN-VASP} = \frac{N_{\text{UN}}}{N_{\text{gen}}} \frac{n_{\text{stable}}}{2,000}.$$

SSUN-VASP uses the same expression with the numerator additionally restricted to non-$P1$ VASP endpoints detected at `symprec=` 0.1 Å. The unique-and-novel condition is the CHGNet pool selection; VASP supplies the stability and final-symmetry tests.

G.5 Aggregation and timing

Each reported DynaCrys configuration and SGEquiDiff is evaluated over five generation seeds; for each of the other main-table baselines, we evaluate their publicly available sample set. We average the per-run statistics without pooling structures, including the per-run medians of E_{hull}. Appendix I reports the corresponding means and sample standard deviations. The VASP window is a single evaluation per row and is not assigned a sampling-seed standard deviation.

Timing covers complete generation through the as-generated structure, including both symbolic and geometric generation for DynaCrys. Generation rate is the number of complete structures per second. All measurements are performed on the same machine equipped with an NVIDIA H100 GPU (80 GB) and 32 CPU processes.

G.6 Geometry sampling steps

We compare 1,000, 500, 250, and 100 geometry sampling steps for DynaCrys on MP-20, using the same symbolic samples and geometry model across schedules. Symbolic sampling retains 200 noise levels. Results are means over five paired generation seeds. All schedules use a fixed budget of 10,000 candidates per seed, including unsuccessful structure construction. The mean SUN@0.1 rates at 1,000 and 500 steps are 41.49% and 38.93%, respectively, so halving the geometry sampling steps retains approximately 93.8% of the discovery rate. Figure 5 shows the complete comparison.

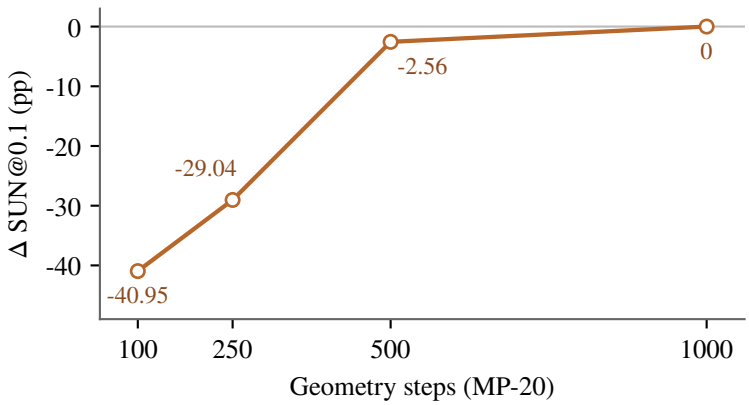


Figure 5: Geometry sampling steps on MP-20. Changes in SUN@0.1 relative to 1,000 geometry steps are averaged over five paired generation seeds of DynaCrys. Rates use the 10,000-candidate budget as the denominator; the symbolic samples and geometry model are held fixed.

H INDEPENDENT EVALUATION WITH MACE-MP-0

We independently relax the MP-20 and MPTS-52 generated sets with MACE-MP-0. For each dataset, we repeat the discovery analysis using the reference data and definitions in Appendix G. The symmetry test uses `symprec`$= 0.1$ Å and 5° angular tolerance on the MACE-MP-0 endpoints. Both DynaCrys configurations lead the compared symmetry-aware baselines across the reported discovery metrics.

Table 11: MP-20 discovery under MACE-MP-0 (%; higher is better). Both DynaCrys configurations and SGEquiDiff show mean ± sample standard deviation over five generation seeds; for the other methods, we evaluate publicly available generated samples. Thresholds @0 and @0.1 are in eV/atom. All values in this table use the MACE-MP-0 evaluation scale.

Method	SUN@0	SSUN@0	SUN@0.1	SSUN@0.1
DiffCSP++	7.93	7.87	34.69	34.40
SymmCD	7.88	7.79	30.04	29.69
WyFormer	6.43	6.40	29.66	29.09
SGEquiDiff	9.11 ± 0.23	8.98 ± 0.24	35.28 ± 0.27	34.51 ± 0.28
DynaCrys (ours)	9.55 ± 0.36	9.44 ± 0.37	38.17 ± 0.32	37.24 ± 0.33
DynaCrys-II (ours)	9.61 ± 0.45	9.53 ± 0.45	38.69 ± 0.53	37.67 ± 0.62

Tables 12 and 13 report discovery and quality results on MPTS-52 under MACE-MP-0. Quality metrics follow the definitions and reference sets in Appendix G, using MACE-MP-0-relaxed endpoints.

Table 12: MPTS-52 discovery under MACE-MP-0 (%; higher is better). Reporting conventions follow Table 11.

Method	SUN@0	SSUN@0	SUN@0.1	SSUN@0.1
SymmCD	6.21	6.07	19.61	19.36
WyFormer	2.98	2.84	9.22	8.95
SGEquiDiff	5.46 ± 0.19	5.32 ± 0.19	22.33 ± 0.27	21.83 ± 0.24
DynaCrys (ours)	7.48 ± 0.24	7.40 ± 0.25	30.22 ± 0.49	29.87 ± 0.47
DynaCrys-II (ours)	6.79 ± 0.18	6.73 ± 0.18	28.52 ± 0.42	28.20 ± 0.45

Table 13: MPTS-52 validity and distributional fidelity under MACE-MP-0. Distributional distances use valid relaxed structures and the training reference. Reporting conventions follow Table 12.

Metric	SymmCD	WyFormer	SGEquiDiff	DynaCrys	DynaCrys-II
Struct. val. (%) ↑	87.150	84.360	97.764 ± 0.218	99.998 ± 0.004	99.998 ± 0.004
Comp. val. (%) ↑	69.94	64.17	78.32 ± 0.21	75.03 ± 0.33	76.93 ± 0.23
d_ρ (g/cm^3) ↓	0.1715	0.7190	0.2598 ± 0.0193	0.2809 ± 0.0133	0.1899 ± 0.0127
$d_{N_{\text{el}}}$ ↓	0.5629	0.2398	0.1482 ± 0.0038	0.0504 ± 0.0038	0.0201 ± 0.0081
d_{Wyckoff} (bits) ↓	0.0155	0.0698	0.0545 ± 0.0015	0.0252 ± 0.0023	0.0168 ± 0.0012
Cov. R (%) ↑	99.80	99.89	99.96 ± 0.02	99.85 ± 0.05	99.85 ± 0.04
Cov. P (%) ↑	85.20	81.20	94.06 ± 0.27	98.56 ± 0.16	98.32 ± 0.10
$P1$ (%) ↓	1.89	4.91	3.79 ± 0.17	1.92 ± 0.08	2.06 ± 0.14
Median E_{hull} (eV/atom) ↓	0.225	0.395	0.185 ± 0.002	0.108 ± 0.002	0.114 ± 0.003

I ACROSS-RUN VARIABILITY

For CHGNet evaluation, we report means and sample standard deviations across five generation seeds. Results for the remaining methods and VASP evaluations are given in the main-text tables.

Table 14: Sampling variability on MP-20 and MPTS-52 (mean ± sample standard deviation). Metrics and units follow Tables 2–6.

(a) Material discovery

Method	SUN@0 (%)	SSUN@0 (%)	SUN@0.1 (%)	SSUN@0.1 (%)
MP-20				
SGEquiDiff	8.45 ± 0.34	8.31 ± 0.35	38.00 ± 0.32	36.38 ± 0.36
DynaCrys (ours)	9.39 ± 0.14	9.31 ± 0.12	41.49 ± 0.34	40.12 ± 0.25
DynaCrys-II (ours)	9.18 ± 0.30	9.10 ± 0.29	41.32 ± 0.51	39.87 ± 0.52
MPTS-52				
SGEquiDiff	5.17 ± 0.20	4.94 ± 0.18	26.29 ± 0.36	24.33 ± 0.35
DynaCrys (ours)	7.30 ± 0.42	7.16 ± 0.44	34.20 ± 0.53	33.05 ± 0.48
DynaCrys-II (ours)	6.65 ± 0.19	6.58 ± 0.19	31.85 ± 0.34	30.85 ± 0.39

(b) Validity and distributional fidelity

Method	Struct. (%)	Comp. (%)	d_ρ	$d_{N_{\mathrm{el}}}$	d_{Wyckoff}
MP-20					
SGEquiDiff	99.50 ± 0.10	86.13 ± 0.43	0.1761 ± 0.0160	0.1195 ± 0.0102	0.0781 ± 0.0029
DynaCrys (ours)	99.97 ± 0.01	83.36 ± 0.19	0.1025 ± 0.0245	0.0518 ± 0.0085	0.0340 ± 0.0032
DynaCrys-II (ours)	99.95 ± 0.02	83.30 ± 0.54	0.1225 ± 0.0358	0.0627 ± 0.0052	0.0344 ± 0.0012
MPTS-52					
SGEquiDiff	98.36 ± 0.16	79.91 ± 0.17	0.2400 ± 0.0228	0.1479 ± 0.0031	0.1266 ± 0.0023
DynaCrys (ours)	99.91 ± 0.01	75.02 ± 0.32	0.2618 ± 0.0149	0.0501 ± 0.0037	0.0626 ± 0.0024
DynaCrys-II (ours)	99.90 ± 0.06	76.90 ± 0.24	0.1779 ± 0.0130	0.0202 ± 0.0083	0.0501 ± 0.0033

Validity and distributional fidelity (continued)

Method	Cov. R (%)	Cov. P (%)	P1 (%)	Median E_{hull} (eV/atom)
MP-20				
SGEquiDiff	99.76 ± 0.04	99.14 ± 0.06	7.95 ± 0.12	0.082 ± 0.001
DynaCrys (ours)	99.71 ± 0.05	99.31 ± 0.10	5.54 ± 0.29	0.072 ± 0.001
DynaCrys-II (ours)	99.70 ± 0.07	99.42 ± 0.07	5.42 ± 0.12	0.060 ± 0.003
MPTS-52				
SGEquiDiff	99.94 ± 0.02	96.24 ± 0.27	10.09 ± 0.27	0.161 ± 0.001
DynaCrys (ours)	99.83 ± 0.07	98.54 ± 0.13	4.59 ± 0.13	0.093 ± 0.003
DynaCrys-II (ours)	99.84 ± 0.05	98.35 ± 0.06	4.66 ± 0.20	0.101 ± 0.002

(c) Geometric quality

	MP-20		MPTS-52	
Method	RMSD (Å)	Matched (%)	RMSD (Å)	Matched (%)
SGEquiDiff	0.190 ± 0.002	91.38 ± 0.10	0.294 ± 0.003	88.05 ± 0.29
DynaCrys (ours)	0.172 ± 0.003	96.18 ± 0.21	0.255 ± 0.003	97.11 ± 0.10
DynaCrys-II (ours)	0.164 ± 0.003	96.65 ± 0.25	0.239 ± 0.003	96.95 ± 0.22

I.1 DETAILED ABLATION RESULTS

Table 15: Detailed ablation results on MPTS-52 under CHGNet, reported as mean ± sample standard deviation over five generation seeds. Fixed G, all-to-all, and w/o FiLM are compared with DynaCrys; MAP is compared with w/o FiLM using the same checkpoint. Discovery and stability rates use the candidate budget; distributional distances use valid relaxed structures and the training reference.

Metric	DynaCrys	Fixed G	All-to-all	w/o FiLM	MAP
Discovery and stability					
SUN@0 (%) ↑	7.30 ± 0.42	6.32 ± 0.25	6.71 ± 0.38	6.51 ± 0.21	5.35 ± 0.19
SSUN@0 (%) ↑	7.16 ± 0.44	6.21 ± 0.28	6.64 ± 0.38	6.41 ± 0.23	5.29 ± 0.18
SUN@0.1 (%) ↑	34.20 ± 0.53	29.43 ± 0.38	31.61 ± 0.51	32.91 ± 0.48	27.65 ± 0.65
SSUN@0.1 (%) ↑	33.05 ± 0.48	28.32 ± 0.40	30.90 ± 0.50	31.73 ± 0.41	27.06 ± 0.64
Stable@0.1 (%) ↑	51.87 ± 0.73	46.00 ± 0.25	50.68 ± 0.56	49.29 ± 0.78	46.73 ± 0.42
Median E_{hull} (eV/atom) ↓	0.093 ± 0.003	0.118 ± 0.001	0.097 ± 0.002	0.102 ± 0.003	0.114 ± 0.002
Validity and coverage					
Struct. val. (%) ↑	99.91 ± 0.01	99.88 ± 0.03	99.94 ± 0.02	99.90 ± 0.03	99.96 ± 0.03
Comp. val. (%) ↑	75.02 ± 0.32	75.81 ± 0.43	74.88 ± 0.26	76.48 ± 0.20	79.74 ± 0.36
Cov. R (%) ↑	99.83 ± 0.07	99.91 ± 0.02	99.85 ± 0.05	99.81 ± 0.05	99.91 ± 0.02
Cov. P (%) ↑	98.54 ± 0.13	97.30 ± 0.25	98.44 ± 0.06	98.49 ± 0.12	97.51 ± 0.17
Distributional and geometric quality					
d_ρ (g/cm^3) ↓	0.2618 ± 0.0149	0.1631 ± 0.0243	0.4148 ± 0.0474	0.4053 ± 0.0255	0.3927 ± 0.0493
$d_{N_{\mathrm{el}}}$ ↓	0.0501 ± 0.0037	0.1120 ± 0.0077	0.0685 ± 0.0041	0.0367 ± 0.0065	0.1386 ± 0.0106
d_{Wyckoff} (bits) ↓	0.0626 ± 0.0024	0.0495 ± 0.0026	0.0573 ± 0.0023	0.0718 ± 0.0021	0.0500 ± 0.0035
RMSD (Å) ↓	0.255 ± 0.003	0.226 ± 0.003	0.247 ± 0.001	0.262 ± 0.001	0.233 ± 0.003

J GENERATED CRYSTAL EXAMPLES

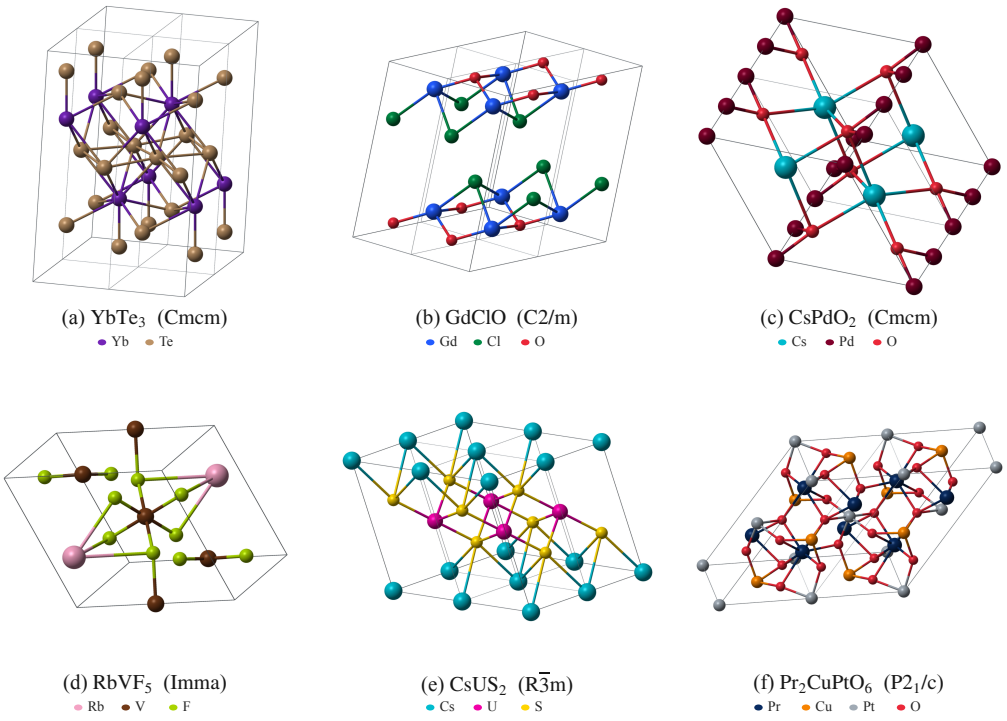


Figure 6: Example crystals generated by DynaCrys on MP-20, shown before relaxation, with their chemical formulae and space groups. Colors distinguish elements.