CRYSTALDLM: FEEDBACK LEARNING OF PERIODIC DIFFUSION LANGUAGE MODELS FOR CRYSTAL GENERATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Crystal generation requires compatible periodic arrangements of atoms, yet geometric validity alone does not guarantee physical quality. We introduce CrystalDLM, a framework built on pretrained masked diffusion language models whose conditional prediction supports both crystal construction and reference-conditioned reconstruction. A lattice-conditioned periodic head couples coordinate candidates in a tractable joint distribution, making periodic compatibility an intrinsic part of draft construction. After diffusion refinement, structured records linking reference structures, executed modifications, and their physical outcomes train a reconstruction diffusion language model to propose targeted replacements and a relative verifier to assess candidates against their references. This feedback stage turns physical evaluations into reusable supervision for subsequent structural decisions. Experiments on MP-20, Perov-5, and MPTS-52 demonstrate higher draft structural validity than existing language-model constructors. On MP-20, the complete framework achieves a stable, unique, and novel (S.U.N.) rate of 10.50%, with component comparisons showing that periodic construction improves structural validity under a common diffusion refiner while feedback learning further increases discovery yield.[1]

1 INTRODUCTION

Discovering novel crystalline materials with targeted properties is critical to accelerating advances in clean energy, photovoltaics, and semiconductor electronics. The search space combines chemical compositions with continuous choices of unit-cell geometry and atomic coordinates, making exhaustive exploration costly (Jain et al., 2013). High-throughput screening and trial-and-error experiments therefore face substantial computational or experimental demands when exploring this space Zeni et al. (2025). Generative models offer a way to focus candidate exploration by learning structural distributions (Xie et al., 2022; Sriram et al., 2024) and proposing crystals for subsequent physical evaluation (Zeni et al., 2025).

Geometric generators model crystal geometry and symmetries: CDVAE, DiffCSP, and MatterGen use diffusion (Xie et al., 2022; Jiao et al., 2023; Zeni et al., 2025), DiffCSP++ adds space-group constraints (Jiao et al., 2024), and FlowMM uses Riemannian flow matching (Miller et al., 2024). Crystal-adapted LLaMA-2 models and CrystaLLM instead learn serialized crystal descriptions with flexible conditioning (Gruver et al., 2024; Antunes et al., 2024). Hybrid methods connect language generation to continuous geometry: CrysLLMGen refines language-generated structures with diffusion, while FlowLLM uses a language model as a learned base distribution for flow matching (Khastagir et al., 2025; Sriram et al., 2024). Nevertheless, these generative approaches generally follow a feedforward design: discrete coordinate drafting typically lacks explicit periodic interactions across cell boundaries, and candidate structures are produced without feedback from downstream physical evaluations. Physical evaluation also guides materials discovery. GenMS uses property predictions in tree search (Yang et al., 2024), while physical evaluations inform reward-based post-training of crystal generators (Park & Walsh, 2025; Wu et al., 2026). Together, these developments motivate learning structural decisions from both the periodic context of a crystal and

[1] Code is available at https://anonymous.4open.science/r/CrystalDLM-E40E.

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.

the physical consequences of earlier modifications. Figure 1 illustrates two aspects of this problem: where an atom can fit geometrically, and which of the feasible arrangements leads to a favorable physical outcome.

First, coordinate choices must be compatible within a periodic structure. Atoms on opposite sides of a unit cell can clash through periodic images (Figure 1, left), while the lattice determines the physical distance corresponding to a fractional displacement. Individually plausible coordinates can therefore form an incompatible arrangement. Periodic relations must influence how unresolved coordinate candidates are combined, so that their compatibility enters the construction decision itself.

Second, choosing among geometrically valid sites can lead to different physical outcomes (Figure 1, right). Evaluating a modification reveals how the resulting structure compares with its starting geometry under a physical criterion. Associating that outcome with the starting structure and the change that produced it creates a learning signal for geometric decisions. The challenge is to make this experience useful beyond the evaluated candidate, by training models to propose and assess modifications to later structures.

Masked diffusion language models (DLMs) recover missing tokens conditioned on partially observed sequences (Sahoo et al., 2024; Nie et al., 2025). Recent work extends them with explicit token dependencies (Li et al., 2026) and reward-informed post-training (Fan et al., 2026; Xie et al., 2026). Their conditional prediction provides a common basis for constructing geometry from an incomplete crystal and predicting replacements conditioned on an existing structure.

To address the above challenges, we introduce **CrystalDLM** to guide structural decisions with periodic context and physical experience. Built on pretrained masked DLMs, it couples coordinate candidates through lattice-conditioned periodic interactions, making their compatibility part of construction. After diffusion refinement, records linking reference structures, executed modifications, and physical outcomes train reference-conditioned reconstruction and relative assessment. Physical experience thus supplies reusable supervision for subsequent proposals and selection. Experiments on three crystal benchmarks show higher draft structural validity than the compared LLM constructors; on MP-20, periodic construction improves structural validity under a common diffusion model, and feedback learning further increases stable or metastable, unique, and novel discovery yields.

2 Related Work

Geometric and language-based crystal generation. CDVAE combines variational modeling with a diffusion decoder, while DiffCSP and MatterGen use diffusion to model crystal structures (Xie et al., 2022; Jiao et al., 2023; Zeni et al., 2025). DiffCSP++ constrains lattice and fractional-coordinate generation by space groups, and FlowMM adapts Riemannian flow matching to crystal geometry and symmetries (Jiao et al., 2024; Miller et al., 2024). Language models provide a complementary representation: CrystaLLM generates crystallographic information files autoregressively, while crystal-adapted LLaMA-2 models support generation and infilling of serialized structures, including the 7B variant evaluated here (Antunes et al., 2024; Gruver et al., 2024). Hybrid methods connect language-generated structures to continuous geometry. CrysLLMGen retains predicted atom types while diffusion refines coordinates and lattice (Khastagir et al., 2025); FlowLLM uses a learned language-model distribution as the starting distribution for Riemannian flow matching (Sriram et al., 2024). FlowLLM is a related hybrid approach beyond our experimental baselines. CrystalDLM builds on construction followed by refinement, with periodic dependencies introduced into coordinate prediction before continuous updates.

Structured prediction in diffusion language models. MDLM and LLaDA recover missing tokens conditioned on partially observed sequences (Sahoo et al., 2024; Nie et al., 2025). This interface can support crystal completion by retaining observed composition and geometry while predicting unfinished fields. However, shared context does not itself couple unresolved coordinate candidates in a factorized output distribution at a fixed decoding step. CoDD introduces explicit token dependencies in diffusion language models (Li et al., 2026), providing a basis for structured predictions beyond separate field scores. CrystalDLM specializes these dependencies to periodic geometry through interactions conditioned on the generated lattice. The same masked-prediction interface also supports reconstruction, where the complete reference and partially modified candidate condition replacement fields.

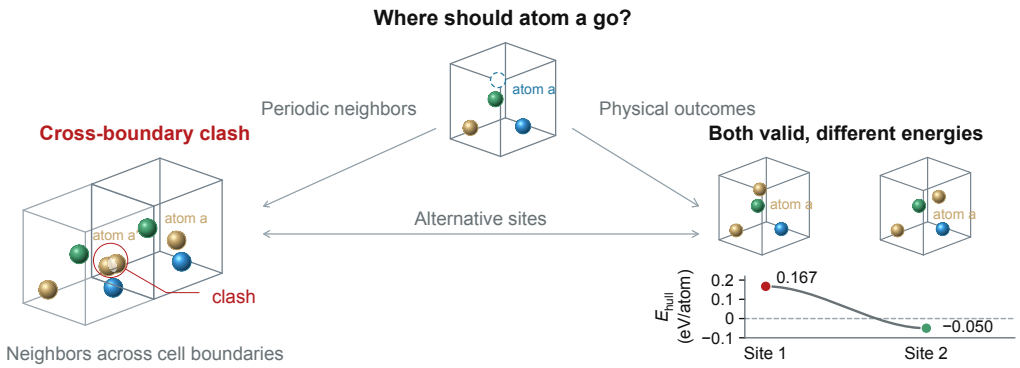

Figure 1: **Periodic context and physical outcomes in crystal generation.** Periodic images reveal a cross-boundary clash (left); different valid sites for the same atom yield different E_{hull}.

Iterative refinement and physical feedback. Thinking with Looped Flows uses temporally aligned denoising objectives to train recurrent states, motivating the reuse of information from earlier computation (Suleymanzade et al., 2026). In materials discovery, GenMS uses property predictions to guide tree search (Yang et al., 2024), while physical evaluations also provide rewards for post-training crystal generators (Park & Walsh, 2025; Wu et al., 2026). Reward-informed post-training has likewise been studied for diffusion language models (Fan et al., 2026; Xie et al., 2026). CrystalDLM uses offline records that associate a reference structure with executed modifications and their physical outcomes. Outcome-selected replacements supervise proposals, while paired outcomes train relative assessment against the unchanged reference.

3 Preliminaries

3.1 Crystal Structure Representation

We represent an N-atom unit cell by $C = (A, X, L)$ (Khastagir et al., 2025; Sriram et al., 2024), where $A = (a_1, \ldots, a_N)$ specifies discrete atomic species, $X \in [0, 1)^{N\times 3}$ contains continuous fractional coordinates, and the rows of a nonsingular $L \in \mathbb{R}^{3\times 3}$ are lattice vectors. Row x_i of X gives Cartesian position $x_i L$. This finite representation specifies the infinite periodic arrangement

$$\widehat{C} = \big\{(a_i, (x_i + n)L) \mid i = 1, \ldots, N,\ n \in \mathbb{Z}^3\big\}. \tag{1}$$

Here, n indexes translated copies of the cell.

3.2 Symmetry and Periodicity

Atom permutations, global translations, joint rotations of Cartesian coordinates and lattice, and alternative unit cells can describe the same crystal (Khastagir et al., 2025; Sriram et al., 2024). Space-group operations preserve its infinite atomic arrangement and species (Levy et al., 2025). In fractional coordinates, periodicity identifies x_i and $x_i + n$ for $n \in \mathbb{Z}^3$. Thus, $[0, 1)^3$ represents a periodic domain whose opposite faces are identified. The shortest separation between distinct sites accounts for all periodic images:

$$d_{ij}(C) = \min_{n\in\mathbb{Z}^3} \big\|(x_i - x_j + n)L\big\|_2, \qquad i \neq j. \tag{2}$$

For any fractional displacement row vector δ, the lattice metric $G_L = LL^\top$ gives $\|\delta L\|_2^2 = \delta G_L \delta^\top$. Thus, identical fractional displacements can have different Cartesian lengths, motivating interactions that respect periodic wrapping and condition on lattice geometry.

3.3 Problem Formulation

Given training crystals $\mathcal{D} = \{C^{(n)}\}_{n=1}^{M}$, de novo crystal generation aims to learn a distribution from which new crystal structures can be sampled:

$$C = (A, X, L) \sim p_\Theta(\,\cdot\,), \tag{3}$$

where Θ denotes the generator parameters. Generation determines species and counts in A, coordinates X, and lattice L, seeking geometric validity and favorable physical outcomes while retaining distributional coverage.

4 METHOD

4.1 OVERVIEW

CrystalDLM connects construction and reconstruction through a refined reference structure (Figure 2). A planner supplies composition A, a periodic DLM completes (X, L), and diffusion refines the draft (Section 4.2). The feedback stage uses offline evaluations to train a reconstruction DLM and a relative verifier. At generation, they propose changes to (X, L) at fixed A and select among the executed candidates and unchanged reference (Section 4.3).

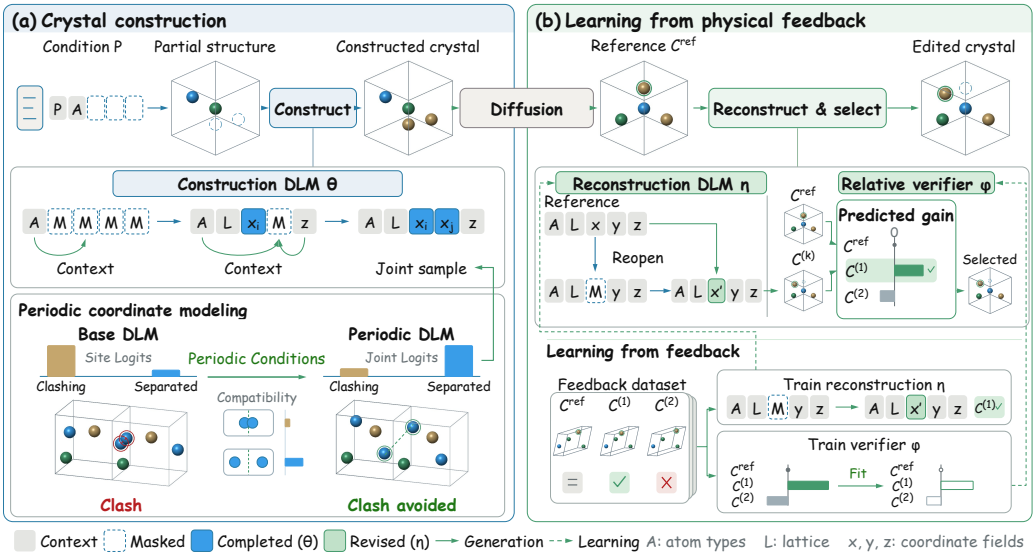

Figure 2: **CrystalDLM framework.** (a) Periodic DLM construction and continuous diffusion produce a reference structure. (b) Physical feedback trains reference-conditioned reconstruction and relative verification to propose and select modifications to this structure.

4.2 PERIODIC CRYSTAL CONSTRUCTION AND REFINEMENT

4.2.1 COMPOSITION AND CONTEXT PLANNING

The planner determines which atoms the constructor will arrange and provides coarse context for their geometry. It samples a condition sequence P from a generation prompt c_{gen}:

$$P \sim p_\omega(P \mid c_{\text{gen}}), \qquad p_\omega(P \mid c_{\text{gen}}) = \prod_{\ell=1}^{|P|} p_\omega(P_\ell \mid P_{<\ell}, c_{\text{gen}}). \tag{4}$$

Autoregressive training fits ω to these sequences. The parsed composition fixes A and N, while descriptors such as lattice category, space group, and volume provide conditioning information. These descriptors do not enforce space-group constraints; the constructor generates the numerical lattice parameters and coordinates. The planner remains fixed during subsequent feedback learning, separating composition sampling from the learned geometric modifications. Appendix A gives the complete condition format and training details.

4.2.2 PERIODIC COORDINATE CONSTRUCTION

Crystal representation and masked prediction. We adapt pretrained LLaDA-8B-Instruct (Nie et al., 2025) to crystal sequences $y = \mathcal{Q}(C)$ containing atom count, six lattice fields for the lengths

and angles of L, and species and fractional-coordinate fields for each atom. Typed tokens represent fields, with continuous values quantized for categorical prediction. Given condition P and a partially observed sequence $\tilde{y}$, the base distribution over masked fields $\mathcal{M}$ is

$$p_\theta^{\mathrm{den}}(y_{\mathcal{M}} \mid P, \tilde{y}) = \prod_{j \in \mathcal{M}} p_{\theta,j}(y_j \mid P, \tilde{y}). \tag{5}$$

The factors share $(P, \tilde{y})$, which updates as completed fields become visible. Adaptation maximizes masked-field likelihood over species and geometry. At generation, the planner fixes A and N, and the constructor completes the lattice fields followed by X. Appendix A gives the serialization, quantization, and adaptation objective.

Lattice-conditioned periodic interactions. At a fixed context, Eq. (5) scores unresolved fields separately, although their coordinates must be compatible under the lattice, including across periodic boundaries. Writing $s = (P, \tilde{y})$, a periodic head couples candidate coordinates using the generated lattice L, species A, and visible geometry. A fixed site-index tree $\mathcal{T}$ defines the pairwise factorization, enabling exact message passing for normalization and sampling at each fixed context. Along the coordinate axis being completed, site i takes a grid value $z_i \in \mathbb{Z}_Q = \{0, \dots, Q-1\}$, representing fractional coordinate z_i/Q. For an edge $(i,j) \in \mathcal{T}$, we define

$$g_{\psi,ij}(\delta; s) = B \tanh\big(a_{\psi,ij}(s)^\top \varphi(\delta)\big), \qquad \delta = z_i - z_j, \tag{6}$$

where $B > 0$ bounds the interaction and φ contains periodic sine and cosine features satisfying $\varphi(\delta + Q) = \varphi(\delta)$. The coefficients $a_{\psi,ij}$ depend on DLM hidden states, species, the lattice metric $G_L = LL^\top$, and visible geometric context. The periodic features respect coordinate wrapping, while lattice conditioning allows the same fractional displacement to receive different compatibility scores under different lattices.

We form a structured candidate distribution (Li et al., 2026) by combining these periodic interactions with the processed unary scores $\ell_{\theta,i}$ from the DLM:

$$\begin{aligned} \Phi_{\theta,\psi}(z; s) &= \sum_i \ell_{\theta,i}(z_i; s) + \sum_{(i,j)\in\mathcal{T}} g_{\psi,ij}(z_i - z_j; s), \\ q_{\theta,\psi}(z \mid s) &= \frac{\exp(\Phi_{\theta,\psi}(z; s)/\tau)}{\sum_{z' \in \Omega_s} \exp(\Phi_{\theta,\psi}(z'; s)/\tau)}, \quad z \in \Omega_s. \end{aligned} \tag{7}$$

Here $\tau > 0$ is the temperature and $\Omega_s = \prod_i \mathcal{S}_i(s)$ is the Cartesian product of per-site supports, with each observed coordinate fixed to its visible value; probability is zero outside this support. The unary scores, interactions, and combined scores correspond to the site logits, compatibility, and joint logits in Figure 2(a).

Learning periodic interactions. We train the periodic head on observed coordinate combinations, keeping the adapted constructor θ frozen. For the true coordinate assignment z^* on this axis, we fit

$$\mathcal{L}_{\mathrm{per}}(\psi) = \mathbb{E}_{C,s}[-\log q_{\theta,\psi}(z^* \mid s)]. \tag{8}$$

The objective fits ψ to observed combinations conditioned on A, L, and visible geometry, learning how periodic compatibility should modify the constructor's unary preferences. Construction then samples coordinates using $q_{\theta,\psi}$. Appendix B gives the feature construction, training views, tree inference, and decoding settings.

4.2.3 Diffusion Geometric Refinement

The constructor predicts lattice and coordinate values on finite grids. To permit continuous geometric updates, we follow CrysLLMGen (Khastagir et al., 2025) and use a separately trained diffusion model to refine the draft at fixed A. Its parameters γ are fitted to coordinate-score and lattice-denoising targets on training crystals. Writing $G = (X, L)$, refinement starts from the constructed geometry G^{init} at intermediate step $t_\star$:

$$G_{t_\star} = G^{\mathrm{init}}, \qquad G_{t-1} \sim r_{\gamma,t}(\cdot \mid G_t, A), \qquad C^{\mathrm{ref}} = (A, X_0, L_0). \tag{9}$$

Here $t = t_\star, \dots, 1$, and $r_{\gamma,t}$ is the reverse transition. The resulting C^{ref} supplies complete coordinates and lattice parameters as the starting geometry for subsequent modifications. Its denoising model is trained on crystal data; the following stage learns from the physical outcomes of changes to these refined references. Appendix A gives training and sampling details.

4.3 PHYSICAL FEEDBACK FOR CRYSTAL RECONSTRUCTION

The feedback stage learns from the physical consequences of modifying refined structures (Figure 2(b)). Its two decisions concern which geometric changes to propose and which resulting candidate to retain. A reconstruction DLM η and verifier ϕ are fitted using records collected offline, then applied to new references during generation. Keeping the starting structure available in both operations allows replacement prediction to use its geometry and selection to compare proposed changes with retaining it.

4.3.1 REFERENCE-CONDITIONED RECONSTRUCTION AND SELECTION

Reference-conditioned reconstruction. The *Reopen* operation masks fields in a candidate copy while retaining the complete reference view $\mathcal{Q}(C^{\mathrm{ref}})$, so predictions can use both the starting geometry and evolving candidate. A modification scope a is formed from learned mode, count, and site predictions under fixed candidate-generation rules. Coordinate scopes replace selected entries of X while preserving L; the full-cell scope permits changes to both X and the six lattice parameters. All scopes preserve atomic species and counts. For unresolved fields $\mathcal{M}(\tilde{y}, a)$ in the candidate $\tilde{y}$, the DLM predicts replacements conditioned on the reference and scope:

$$\hat{y}_j \sim p_{\eta,j}\big(\cdot \mid P, \mathcal{Q}(C^{\mathrm{ref}}), \tilde{y}, a\big), \qquad j \in \mathcal{M}(\tilde{y}, a). \tag{10}$$

Iterative masked completion updates $\tilde{y}$ as predictions are committed; executing the proposal produces $C^{(k)}$, with unchanged numerical fields retaining their continuous reference values. Appendix C specifies scope generation and patch execution.

Selecting a candidate relative to its reference. For each executed candidate $C^{(k)}$, h_k contains frozen reconstruction features and geometric changes relative to C^{ref}. The verifier assigns a physical-quality score $f_\phi(h_k)$. Using unchanged-reference features h_0, its *Predicted gain* is

$$\widehat{g}_k = f_\phi(h_k) - f_\phi(h_0). \tag{11}$$

Base selection returns the highest-scoring executable candidate or the unchanged reference, whose zero gain allows it to be retained when no proposed change has positive predicted gain. The complete selection rule and bounded optional field revision are specified in Appendices C.2 and D.2.

4.3.2 LEARNING FROM PHYSICAL FEEDBACK

Evaluated modifications. Each record b_k links the reference C^{ref}, executed replacement fields, submitted candidate $C^{(k)}$, and outcomes obtained under the same relaxation and screening protocol. Labels combining novelty with stability or metastability define a normalized utility $u(C)$ and a relative gain $\Delta u_k = u(C^{(k)}) - u(C^{\mathrm{ref}})$ (Appendix C). Replacement fields supply proposal targets; paired outcomes supervise their relative assessment. Grouping records by reference places alternatives on a common physical comparison and preserves the connection between each change and its observed outcome.

Learning what to propose. We first fit the reconstruction model to modifications selected by an improvement rule based on physical quality and energy above the hull (Appendix C). A subsequent refit uses differences in measured utility to emphasize more favorable alternatives for the same reference. For an ordered pool $\mathcal{B} = (b_0, \ldots, b_K)$ of alternatives for one reference, including the unchanged structure b_0, a finite-pool teacher assigns weights

$$q_T(k \mid \mathcal{B}) = \frac{\exp(\Delta u_k/\beta)}{\sum_{l=0}^{K} \exp(\Delta u_l/\beta)}, \qquad \beta > 0. \tag{12}$$

Higher-utility alternatives receive more weight. For eligible references, teacher-sampled records are mixed with the initial supervision (Appendix C); Appendix C.1 gives the optimization interpretation. The target $y^\star$ contains the proposal values actually used to form the submitted candidate; relaxation outcomes select and weight these targets. Selected records and random partial views of their target fields define $\mathcal{D}_{\mathrm{fb}}$. With conditioning ν from Eq. (10), we supervise the remaining fields $\mathcal{M}_\nu$:

$$\mathcal{L}_{\mathrm{rec}}(\eta) = \mathbb{E}_{\mathcal{D}_{\mathrm{fb}}}\left[-\frac{1}{m_\nu}\sum_{j\in\mathcal{M}_\nu} \log p_{\eta,j}(y_j^\star \mid \nu)\right] + \lambda \mathcal{R}_{\mathrm{KL}}(\eta), \tag{13}$$

where $m_\nu = \max(1, |\mathcal{M}_\nu|)$ and $\lambda \geq 0$. The first term increases the likelihood of outcome-selected replacements for permitted fields of X and L, while the forward-KL regularizer $\mathcal{R}_{\mathrm{KL}}$ limits departures from a frozen starting reconstruction checkpoint. These updates transfer evaluated experience into later replacement predictions; Appendix C gives decision supervision and the expanded loss.

Learning relative assessments. After fitting the reconstruction DLM, we train the verifier on frozen reconstruction features and geometric changes. Modifications grouped by reference provide paired outcome targets for the quality difference in Eq. (11), teaching selection to compare each candidate with the option of retaining its starting structure. Appendices C.2 and D.1 give the targets, regression objective, and training sequence.

5 Experiments

5.1 Experimental Setup

Datasets. We evaluate construction quality and complete-generation outcomes on MP-20, Perov-5, and MPTS-52. MP-20 and Perov-5 follow the datasets curated by Xie et al. (2022): MP-20 contains inorganic crystals with at most 20 atoms per unit cell from the Materials Project (Jain et al., 2013), while Perov-5 contains five-atom perovskites from computational screening (Castelli et al., 2012). MPTS-52 extends evaluation to cells with at most 52 atoms (Baird et al., 2024). These benchmarks cover a fixed-size crystal family, varied inorganic compositions, and larger unit cells, allowing us to compare construction across different structural settings.

Baselines. We compare with geometric generators CDVAE (Xie et al., 2022), MatterGen (Zeni et al., 2025), DiffCSP (Jiao et al., 2023), DiffCSP++ (Jiao et al., 2024), and FlowMM (Miller et al., 2024); a crystal-adapted LLaMA-2 (7B) (Gruver et al., 2024); and the hybrid CrysLLMGen (Khastagir et al., 2025). CrysLLMGen provides the closest language–diffusion comparison. CDVAE and MatterGen are omitted on MPTS-52 because of reproduction and sampling difficulties in our runs.

Evaluation metrics. Following Xie et al. (2022), *Structural* and *Compositional* validity check geometric admissibility and composition plausibility. Coverage *Precision* measures generated structures' proximity to the reference set, whereas *Recall* measures how much of that set is covered. Property metrics $\mathrm{wdist}(\rho)$ and $\mathrm{wdist}(N_{\mathrm{el}})$ are one-dimensional Wasserstein-1 distances between generated and reference distributions of mass density ρ and distinct-element count N_{el}, respectively; N_{el} differs from atom count N. Smaller distances mean closer distributions. On MP-20, V.U.N., S.U.N., and M.S.U.N. jointly require validity, stability, or metastability, respectively, together with uniqueness within the generated collection and novelty relative to the training-reference set (Mal et al., 2026; Miller et al., 2024; Khastagir et al., 2025). Higher rates are better.

Evaluation collections. Complete-generation and draft comparisons each use 10,000 parseable outputs per dataset and method; failed generations are resampled, and sampling success reports the fraction of all attempts that are parseable. Physical discovery rates use a separate collection of 5,000 outputs per method on MP-20, with the full collection as denominator. Appendix A gives training and sampling settings; Appendix E defines the metrics, matching rules, and evaluation protocol.

5.2 De Novo Crystal Generation

CrystalDLM improves structural validity and Coverage Recall over CrysLLMGen across all three datasets, while Coverage Precision is lower (Table 1). These metrics evaluate final outputs before physical relaxation. The structural-validity difference is largest on MPTS-52, at 99.65% versus 97.96%. On MP-20, Recall rises from 97.13% to 99.44%, indicating broader reference-set coverage under the fingerprint criterion. The accompanying lower Precision means that a smaller fraction of generated structures lies near the reference set, so the coverage comparison reflects different aspects of distributional agreement. Element-count distribution distance also decreases on every dataset relative to CrysLLMGen.

Against geometric generators, CrystalDLM matches the highest structural validity on MP-20 and Perov-5 and the highest Recall on MP-20; leading geometric methods retain higher Recall on the

Table 1: De novo crystal generation on MP-20, Perov-5, and MPTS-52. Throughout the comparison tables, bold and underlined values mark the best and second-best results per dataset and metric, respectively.

Dataset	Model	Validity (%)		Coverage (%)		Property ↓	
		Structural ↑	Compositional ↑	Precision ↑	Recall ↑	wdist(ρ)	wdist(N_{el})
MP-20	CDVAE	<u>99.97</u>	86.72	99.33	99.28	0.72	0.28
	MatterGen	**100.00**	86.87	<u>99.43</u>	99.27	0.46	0.24
	DiffCSP	99.76	83.82	99.40	<u>99.33</u>	<u>0.30</u>	0.52
	DiffCSP++	99.82	84.74	99.34	**99.44**	**0.13**	0.52
	FlowMM	92.70	84.06	99.26	98.64	1.45	<u>0.10</u>
	LLaMA-2 (7B)	87.91	<u>91.76</u>	96.01	91.11	0.65	0.96
	CrysLLMGen	99.95	91.59	**99.84**	97.13	0.39	0.12
	CrystalDLM	**100.00**	**92.31**	98.95	**99.44**	0.79	**0.07**
Perov-5	CDVAE	**100.00**	98.62	98.48	**99.45**	0.13	0.06
	MatterGen	**100.00**	98.23	98.43	<u>98.99</u>	<u>0.11</u>	<u>0.03</u>
	DiffCSP	**100.00**	98.50	98.18	98.84	0.15	0.04
	DiffCSP++	**100.00**	98.50	<u>98.60</u>	98.76	**0.09**	**0.02**
	FlowMM	**100.00**	98.22	89.12	98.34	0.78	0.04
	LLaMA-2 (7B)	99.79	99.29	98.57	93.56	0.65	0.08
	CrysLLMGen	<u>99.98</u>	<u>99.44</u>	**98.73**	94.06	0.21	0.12
	CrystalDLM	**100.00**	**99.83**	<u>98.60</u>	97.08	0.21	0.06
MPTS-52	DiffCSP	<u>99.74</u>	65.46	88.82	<u>99.30</u>	**0.82**	0.49
	DiffCSP++	**99.86**	78.26	<u>97.06</u>	**99.63**	1.03	**0.45**
	FlowMM	93.66	72.80	**98.22**	99.10	1.98	<u>0.47</u>
	LLaMA-2 (7B)	91.25	<u>86.73</u>	96.10	96.32	1.27	0.89
	CrysLLMGen	97.96	85.44	96.38	97.83	1.05	0.86
	CrystalDLM	99.65	**87.22**	95.85	98.81	<u>0.86</u>	**0.45**

other datasets. It also achieves the highest compositional validity among the evaluated methods on all three datasets. Because the geometric refinement and feedback stages preserve atom types, this composition comparison concerns the complete generator rather than an isolated geometric component. Property agreement remains mixed: CrystalDLM achieves the best or joint-best element-count distribution distance on MP-20 and MPTS-52, but its MP-20 density distribution distance is 0.79 versus the best value of 0.13. These results motivate assessing physical discovery yield alongside geometric and distributional metrics.

5.3 Stable and Novel Crystal Generation

Geometric validity alone does not establish energetic stability or novelty. S.U.N. is the fraction of generated structures simultaneously satisfying stability, uniqueness, and novelty; M.S.U.N. replaces stability with metastability. These are intersections of conditions, not products of marginal rates. We screen CHGNet-relaxed energies against a Materials Project competing-phase hull (Deng et al., 2023; Jain et al., 2013), using $E_{\text{hull}} \leq 0$ for stability and $E_{\text{hull}} \leq 0.1$ eV/atom for metastability. The latter includes stable structures. Both rates use the full separate MP-20 collection of 5,000 outputs as denominator; matching rules and evaluation timing are in Appendix E.

CrystalDLM yields more unique, novel candidates meeting either energetic threshold than the compared methods (Table 2). Its S.U.N. rate is 10.50%, versus CrysLLMGen's 9.42%; its M.S.U.N. rate is 51.24%, versus MatterGen's 49.44%. The gains concern discovery yield under the stated CHGNet screening protocol, rather than geometric validity alone.

Table 2: V.U.N., S.U.N., and M.S.U.N. rates on MP-20 (%).

Model	V.U.N. ↑	S.U.N. ↑	M.S.U.N. ↑
CDVAE	73.24	3.72	31.96
MatterGen	77.20	9.10	<u>49.44</u>
DiffCSP	72.52	8.86	48.22
DiffCSP++	74.40	9.28	40.84
CrysLLMGen	<u>77.84</u>	<u>9.42</u>	44.78
CrystalDLM	**81.68**	**10.50**	**51.24**

5.4 Analysis of Model Components

5.4.1 Crystal Construction before Diffusion

Table 3: Quality of constructed crystals before continuous diffusion and feedback reconstruction.

Dataset	Model	Sampling success (%) ↑	Validity (%) Structural ↑	Validity (%) Compositional ↑	Coverage (%) Precision ↑	Coverage (%) Recall ↑	Property ↓ wdist(ρ)	Property ↓ wdist(N_{el})
MP-20	LLaMA-2 (7B)	98.27	87.91	91.76	96.01	**91.11**	0.65	0.96
	CrysLLMGen (draft)	97.85	96.28	89.92	97.46	84.51	**0.42**	0.11
	CrystalDLM (draft)	**99.46**	**99.83**	**92.31**	**98.53**	90.60	1.21	**0.07**
Perov-5	LLaMA-2 (7B)	**99.99**	99.79	99.29	98.57	93.56	0.65	0.08
	CrysLLMGen (draft)	99.97	99.95	99.11	**98.70**	**93.63**	0.32	0.12
	CrystalDLM (draft)	**99.99**	**100.00**	**99.83**	96.43	92.60	**0.14**	**0.06**
MPTS-52	LLaMA-2 (7B)	85.86	91.25	86.73	96.10	**96.32**	1.27	0.89
	CrysLLMGen (draft)	77.39	90.52	85.44	97.10	96.02	1.38	0.86
	CrystalDLM (draft)	**98.22**	**97.13**	**87.22**	**99.00**	94.41	**0.76**	**0.45**

Before continuous refinement, periodic DLM drafts have higher structural validity than the LLaMA-2 and CrysLLMGen drafts on all three datasets (Table 3). Sampling success measures parseability; the other metrics characterize the collected drafts. On MPTS-52, CrystalDLM attains 98.22% sampling success and 97.13% structural validity, versus 77.39% and 90.52% for CrysLLMGen, placing both improvements at the draft stage. Element-count distribution distance is lower than both LLM constructors' on every dataset, although draft Recall remains below LLaMA-2's. The draft comparison therefore supports stronger initial geometric validity without implying uniformly better distributional coverage. The following controlled comparisons examine how periodic interactions and feedback contribute under a shared diffusion model.

5.4.2 Ablation Studies

All variants in Table 4 share the planner, condition list, constructor checkpoint, and diffusion model. DLM uses the base constructor. Periodic DLM adds periodic coordinate interactions. CrystalDLM applies the complete feedback stage to the same refined references as Periodic DLM.

Table 4: Ablation of periodic construction and physical feedback on MP-20 (%).

Model	Structural validity ↑	S.U.N. ↑	M.S.U.N. ↑
DLM	98.40	7.86	45.08
Periodic DLM	**100.00**	8.48	47.24
CrystalDLM	**100.00**	**10.50**	**51.24**

Periodic interactions raise structural validity from 98.40% to 100.00% after the same diffusion stage, with accompanying increases in S.U.N. and M.S.U.N. From the shared references, feedback further increases S.U.N. and M.S.U.N. by 2.02 and 4.00 percentage points. Both variants achieve 100.00% validity in the separate structural evaluation, while their physical discovery rates differ, illustrating the complementary information provided by the two evaluations. The feedback result measures the joint effect of proposals, selection, and associated inference operations; it does not isolate the benefit of each operation. Together, the comparisons connect periodic compatibility to improved construction and the complete feedback stage to higher physical discovery yield.

6 Conclusion

We introduced CrystalDLM to combine periodic construction with learning from physical feedback. Lattice-conditioned interactions coupled coordinate candidates during construction, while evaluated modifications supplied supervision for reference-conditioned reconstruction and relative assessment. Experiments demonstrated stronger draft validity than the compared LLM constructors. On MP-20, periodic construction improved structural validity under a common diffusion model, and the complete feedback stage increased S.U.N. and M.S.U.N. yields from shared references. These findings supported using periodic context and accumulated physical outcomes to guide structural decisions. Future work will explore more accurate physical feedback through higher-fidelity evaluations and extend the framework to generation conditioned on target material properties.

AI USE STATEMENT

Generative AI tools assisted in code development, repository inspection, literature review, research ideation and execution, mathematical proofs, and manuscript drafting and editing. The authors take responsibility for the final content

REPRODUCIBILITY STATEMENT

Appendices A–E describe the model, training and sampling settings, and evaluation protocol. Code and configurations are available in the anonymous repository.

REFERENCES

AI@Meta. Llama 3 model card. Official model card, 2024. URL https://huggingface.co/meta-llama/Meta-Llama-3-8B.

Luis M. Antunes, Keith T. Butler, and Ricardo Grau-Crespo. Crystal structure generation with autoregressive large language modeling. *Nature Communications*, 15:10570, 2024. doi: 10.1038/s41467-024-54639-7. URL https://www.nature.com/articles/s41467-024-54639-7.

Sterling G. Baird, Hasan M. Sayeed, Joseph Montoya, and Taylor D. Sparks. matbench-genmetrics: A Python library for benchmarking crystal structure generative models using time-based splits of Materials Project structures. *Journal of Open Source Software*, 9(97):5618, 2024. doi: 10.21105/joss.05618. URL https://joss.theoj.org/papers/10.21105/joss.05618.

Ivano E. Castelli, David D. Landis, Kristian S. Thygesen, Søren Dahl, Ib Chorkendorff, Thomas F. Jaramillo, and Karsten W. Jacobsen. New cubic perovskites for one- and two-photon water splitting using the computational materials repository. *Energy & Environmental Science*, 5(10):9034–9043, 2012. doi: 10.1039/C2EE22341D. URL https://pubs.rsc.org/en/content/articlehtml/2012/ee/c2ee22341d.

Daniel W. Davies, Keith T. Butler, Adam J. Jackson, Jonathan M. Skelton, Kazuki Morita, and Aron Walsh. SMACT: Semiconducting materials by analogy and chemical theory. *Journal of Open Source Software*, 4(38):1361, 2019. doi: 10.21105/joss.01361. URL https://joss.theoj.org/papers/10.21105/joss.01361.

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. URL https://www.nature.com/articles/s42256-023-00716-3.

Yijia Fan, Jing Yang, Mingyu Liu, Kaitong Cai, Jian Wang, Keze Wang, and Jusheng Zhang. Reinforcement learning for diffusion LLMs via energy-based Gibbs alignment. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 45938–45948, 2026. doi: 10.18653/v1/2026.acl-long.2131. URL https://aclanthology.org/2026.acl-long.2131/.

Nate Gruver, Anuroop Sriram, Andrea Madotto, Andrew Gordon Wilson, C. Lawrence Zitnick, and Zachary Ulissi. Fine-tuned language models generate stable inorganic materials as text. In *International Conference on Learning Representations*, 2024. URL https://proceedings.iclr.cc/paper_files/paper/2024/hash/5dc6d1c81754c32008c9667339b3fdd3-Abstract-Conference.html.

Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=nZeVKeeFYf9.

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. URL https://doi.org/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*, 2023. URL https://arxiv.org/abs/2309.04475.

Rui Jiao, Wenbing Huang, Yu Liu, Deli Zhao, and Yang Liu. Space group constrained crystal generation. In *International Conference on Learning Representations*, 2024. URL https://proceedings.iclr.cc/paper_files/paper/2024/hash/1ba1b2133ac3d3dda0ce5cd67dab8456-Abstract-Conference.html.

Subhojyoti Khastagir, Kishalay Das, Pawan Goyal, Seung-Cheol Lee, Satadeep Bhattacharjee, and Niloy Ganguly. LLM meets diffusion: A hybrid framework for crystal material generation. In *Advances in Neural Information Processing Systems*, 2025. URL https://proceedings.neurips.cc/paper_files/paper/2025/hash/f789a628fca473e922c806657512a20f-Abstract-Conference.html.

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*, 2025. URL https://arxiv.org/abs/2502.03638.

Ian Li, Zilei Shao, Benjie Wang, Rose Yu, Guy Van den Broeck, and Anji Liu. Breaking the factorization barrier in diffusion language models. arXiv:2603.00045, 2026. URL https://arxiv.org/abs/2603.00045.

Sourav Mal, Nehad Ahmed, Junaid Jami, Subhankar Mishra, and Prasenjit Sen. DiffCrysGen: a generative diffusion model for accelerated design of inorganic crystalline materials. *npj Computational Materials*, 12:282, 2026. doi: 10.1038/s41524-026-02147-1. URL https://www.nature.com/articles/s41524-026-02147-1.

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*, volume 235 of *Proceedings of Machine Learning Research*, pp. 35664–35686, 2024. URL https://proceedings.mlr.press/v235/miller24a.html.

Shen Nie, Fengqi Zhu, Zebin You, Xiaolu Zhang, Jingyang Ou, Jun Hu, Jun Zhou, Yankai Lin, Ji-Rong Wen, and Chongxuan Li. Large language diffusion models. In *Advances in Neural Information Processing Systems*, volume 38, 2025. URL https://proceedings.neurips.cc/paper_files/paper/2025/hash/48b383b24230e0e6e649d9c98dae4d8c-Abstract-Conference.html.

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. URL https://doi.org/10.1016/j.commatsci.2012.10.028.

Hyunsoo Park and Aron Walsh. Guiding generative models to uncover diverse and novel crystals via reinforcement learning. arXiv:2511.07158, 2025. URL https://arxiv.org/abs/2511.07158.

Subham Sekhar Sahoo, Marianne Arriola, Yair Schiff, Aaron Gokaslan, Edgar Marroquin, Justin T. Chiu, Alexander Rush, and Volodymyr Kuleshov. Simple and effective masked diffusion language models. In *Advances in Neural Information Processing Systems*, 2024. URL https://arxiv.org/abs/2406.07524.

Anuroop Sriram, Benjamin Kurt Miller, Ricky T. Q. Chen, and Brandon M. Wood. FlowLLM: Flow matching for material generation with large language models as base distributions. In *Advances in Neural Information Processing Systems*, 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/51d317df78eded9eb3c9d3fb1091c279-Abstract-Conference.html.

Ayhan Suleymanzade, Chanhyuk Lee, Floor Eijkelboom, Nicholas M. Boffi, İsmail İlkan Ceylan, and Jinwoo Kim. Thinking with looped flows. arXiv:2609.11801, 2026. URL https://arxiv.org/abs/2609.11801.

Logan Ward, Ankit Agrawal, Alok Choudhary, and Christopher Wolverton. A general-purpose machine learning framework for predicting properties of inorganic materials. *npj Computational Materials*, 2:16028, 2016. doi: 10.1038/npjcompumats.2016.28. URL https://www.nature.com/articles/npjcompumats201628.

Yuyang Wu, Stefano Falletta, Delia McGrath, and Sherry Yang. CrystalReasoner: Reasoning and RL for property-conditioned crystal structure generation. arXiv:2605.14344, 2026. URL https://arxiv.org/abs/2605.14344.

Shaoan Xie, Lingjing Kong, Xiangchen Song, Xinshuai Dong, Guangyi Chen, Eric P. Xing, and Kun Zhang. Advancing reasoning in diffusion language models with denoising process rewards. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 42703–42720, 2026. doi: 10.18653/v1/2026.acl-long.1978. URL https://aclanthology.org/2026.acl-long.1978/.

Tian Xie, Xiang Fu, Octavian-Eugen Ganea, Regina Barzilay, and Tommi Jaakkola. Crystal diffusion variational autoencoder for periodic material generation. In *International Conference on Learning Representations*, 2022. URL https://arxiv.org/abs/2110.06197.

Sherry Yang, Simon Batzner, Ruiqi Gao, Muratahan Aykol, Alexander L. Gaunt, Brendan McMorrow, Danilo J. Rezende, Dale Schuurmans, Igor Mordatch, and Ekin D. Cubuk. Generative hierarchical materials search. In *Advances in Neural Information Processing Systems*, 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/447d012bd95b6767a4bfdebf96cdfcc9-Abstract-Conference.html.

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. URL https://www.nature.com/articles/s41586-025-08628-5.

A Representation and Model Training

Field order and serialization. We serialize an N-atom crystal in the order

$$\big(N, \ell_a, \ell_b, \ell_c, \alpha, \beta, \gamma, (a_i, x_{i1}, x_{i2}, x_{i3})_{i=1}^{N}\big), \qquad T_C = 1 + 6 + 4N,$$

where ℓ_a, ℓ_b, ℓ_c and α, β, γ specify the lattice lengths and angles. Each field maps to one special token: <N_...> for atom count; <LA_...>, <LB_...>, and <LC_...> for lengths; <AA_...>, <AB_...>, and <AG_...> for angles; <E_...> for species; and <X_...>, <Y_...>, and <Z_...> for fractional coordinates. The crystal answer contains exactly N site entries, without empty-site slots or site-index tokens. Its length T_C excludes the condition prompt and batch padding. We use row lattice vectors, so $G_L = LL^\top$ is the fractional-coordinate metric. Sites are grouped in the condition's element order (increasing atomic number for internally formed conditions); parsing order is retained within each group. Data preparation adds no cell reduction. The inherited continuous-model interface later applies Niggli reduction without primitive-cell conversion.

Numerical precision and vocabulary limits. Length tokens use integer bins $0, \ldots, 500$ with spacing 0.1 Å, angle tokens use integer degrees $1, \ldots, 179$, and species tokens cover atomic numbers $1, \ldots, 94$. Fractional coordinates are wrapped into $[0, 1)$ and rounded to the nearest 0.01; ties round upward. Thus the token vocabulary contains coordinate bins $0, \ldots, 100$ for each axis, with bins 0 and 100 representing the same periodic position. The periodic head merges these aliases into $Q = 100$ distinct states (Appendix B). Lengths and angles are rounded and clipped to their configured ranges, with clipping recorded by the encoder. These vocabulary limits are distinct from the geometric legality checks used during generation. The reference dataset configurations allow $1 \le N \le 20$ for MP-20, $N = 5$ for Perov-5, and $1 \le N \le 52$ for the larger-cell benchmark. Their maximum crystal-answer lengths are therefore 87, 27, and 215 tokens, respectively. Each configuration requires a tokenizer and checkpoint with the corresponding atom-count vocabulary.

Encoding and masking example. Consider a two-atom cell with species $(\mathrm{Na}, \mathrm{Cl})$, lengths $(4.14, 5.02, 6.00)$ Å, angles $(90, 90, 90)$ degrees, and fractional coordinates $(0.127, 0, 0)$ and $(0.5, 0.5, 0.5)$. Its 15-field answer is

```
<N_002><LA_041><LB_050><LC_060><AA_090><AB_090><AG_090>
<E_Na><X_013><Y_000><Z_000>
<E_Cl><X_050><Y_050><Z_050>
```

Decoding recovers lengths $(4.1, 5.0, 6.0)$ Å and the first fractional coordinate 0.13. Writing [MASK] for the DLM mask token, hiding only this coordinate changes the first site entry to

<E_Na>[MASK]<Y_000><Z_000>.

The condition and other fields remain visible, and the target for that position is the complete token <X_013>. During construction, the planner fixes count and species in advance, and the model completes the numerical lattice and coordinate fields.

Why this representation fits masked generation. One token per field aligns a masked prediction with an entire species or numerical value. Typed vocabularies make the allowed values for a coordinate, angle, or length explicit, while dynamic site entries keep answer length linear in the actual atom count. Fractional coordinates provide a common periodic domain on which aliases and coordinate differences can be handled consistently. The 0.01 grid is a finite-precision design choice that bounds the coordinate vocabulary and the $O(NQ^2)$ tree-inference cost; its physical resolution depends on the lattice. Subsequent diffusion operates on continuous coordinates and lattice variables. These choices align fieldwise completion with periodic candidate prediction.

Coordinate and lattice changes. For matched sites expressed in the same cell basis, let $x_i' = x_i + \delta x_i$ and $L' = L + \delta L$, with fixed coordinate representatives. Since the Cartesian position is $r_i = x_i L$, the finite change is

$$\begin{aligned} r_i' - r_i &= (x_i + \delta x_i)(L + \delta L) - x_i L \\ &= \delta x_i L + x_i \delta L + \delta x_i \delta L. \end{aligned} \tag{A.1}$$

The terms separate the effects of coordinate changes, lattice changes, and their interaction. A lattice update can therefore change interatomic distances while fractional coordinates remain fixed. Together with the metric $G_L = LL^\top$, this relation explains why coordinate construction conditions on the lattice and why full-cell reconstruction considers both X and L.

Planner. The planner adapts Meta-Llama-3-8B with LoRA to produce composition and coarse conditions from material descriptions (AI@Meta, 2024; Hu et al., 2022). Its structured answer contains formula, anion, charge, lattice category, space group, and volume context. The parsed composition determines the species and counts supplied to the geometric constructor; the remaining descriptors serve as conditioning information, without enforcing space-group symmetry. A supplied condition may replace planner sampling. The sequence law in Eq. (4) is autoregressive over the condition tokens, including its end marker; sampling temperature and truncation settings are implementation parameters. Training uses autoregressive answer prediction in two one-epoch stages, with the implementation defaults in Table 5. The planner is fitted before geometric generation and remains fixed during feedback learning.

Masked adaptation. The geometric DLM initializes from LLaDA-8B-Instruct. Its LoRA factors and full resized input/output vocabulary tables are trained; other pretrained weights are frozen. For each training crystal, draw $t \sim \mathcal{U}[0,1]$ and independently mask answer fields with probability $\rho_t = \epsilon + (1-\epsilon)t$, where $\epsilon = 10^{-3}$. If no answer field is selected, mask one uniformly. With $T_C = 7+4N$, the objective is

$$\mathcal{L}_{\mathrm{adapt}}(\theta) = \mathbb{E}_{C,t,\mathcal{M}_t}\left[-\frac{1}{T_C \rho_t}\sum_{j\in\mathcal{M}_t}\log p_{\theta,j}(y_j \mid P, \tilde{y})\right]. \tag{A.2}$$

The implementation averages the inverse-mask-probability-weighted token losses over each complete answer and then averages over the batch. Species, count, and geometry fields all belong to the training answer; prompt and padding positions receive no loss. The softmax covers the augmented vocabulary, and the vocabulary tables learn the newly added crystal tokens together with the adapted backbone. We use this objective, including the nonempty-mask safeguard, as a masked-denoising surrogate. At generation, count and species are fixed from the condition before numerical geometry is completed. The lattice fields are generated first, followed by coordinate axes and element blocks as specified in Appendix B.

Diffusion refinement. The refiner follows CrysLLMGen and DiffCSP (Khastagir et al., 2025; Jiao et al., 2023). With A fixed, training perturbs lattice matrices and fractional coordinates as

$$L_t = \sqrt{\bar{\alpha}_t}L_0 + \sqrt{1-\bar{\alpha}_t}\epsilon_L, \qquad X_t = \mathrm{wrap}(X_0 + \sigma_t \epsilon_X), \tag{A.3}$$

where ϵ_L, ϵ_X are independent standard Gaussian arrays, $\bar{\alpha}_t = \prod_{s=1}^{t}(1-\beta_s)$, and $\mathrm{wrap}(X) = X - \lfloor X \rfloor$ acts elementwise. The network takes (A, X_t, L_t, t) and is trained independently of the DLMs using the sum of mean squared errors for lattice noise and the normalized negative score of the wrapped coordinate perturbation. At generation, the initial structure is converted to the refiner's cell representation and its geometry is injected directly at an intermediate step $t_\star$, without additional forward corruption. The reverse transitions in Eq. (9) produce C^{ref}; the reference default is $t_\star = 800$ out of 1,000 timesteps. This generative refinement is separate from CHGNet relaxation used for physical labels. Failed or invalid refinement retains the initial structure, and an unencodable continuous reference is retained without reconstruction.

Reconstruction conditioning. The reconstruction copy keeps inherited vocabulary tables fixed. It adds separate residual embeddings from the reference token view, the evolving candidate view, their aligned differences, and task/scope/visibility features. Each geometric view uses species, periodic coordinate features, and finite-image radial messages under its decoded lattice. Unknown coordinates and invalid lattice features are gated out. The reference view supplies conditioning embeddings aligned with the candidate fields. Conditioning uses token-derived geometry; patch execution retains the unquantized reference separately. The four scopes are identity, selected-site coordinates, all coordinates, and lattice plus all coordinates. Local site counts are drawn from feasible values in $\{1, 2, 4, 8\}$; count and species are fixed.

Warm-up and parameter paths. Warm-up corrupts selected numerical fields in training crystals and learns their clean values. It supervises the next field after a randomly revealed prefix, together with scope/count/site labels. Physical-outcome supervision is introduced in the subsequent feedback fit. Numerical reconstruction and forward KL update LoRA and conditioning modules, while decision losses use detached hidden states and update their heads. The teacher-guided refit updates LoRA and decision heads with conditioning modules frozen. The verifier is fitted afterward on the final frozen reconstruction representation. Table 5 summarizes the reference implementation settings.

Table 5: Default training and sampling settings of the reference implementation.

Stage	Settings
Planner	Two one-epoch stages; AdamW 2×10^{-5}; effective batch 8; LoRA rank 16, scale 32.
DLM adaptation	Two epochs; AdamW 5×10^{-5}; effective batch 16; LoRA rank 8, scale 32.
Periodic head	One epoch; Adam 10^{-4}; batch 16; width 64; $H = 8$, $B = 4$; validation NLL checkpoint.
Continuous refiner	500 epochs; Adam 10^{-3}; batch 512; 1,000 diffusion timesteps; final checkpoint.
Reconstruction warm-up	One epoch; AdamW 10^{-5}; batch 16, microbatch 4.
Feedback fit	Eight epochs; batch 16; content LR 2×10^{-6}; head LR 10^{-4}; reference KL weight 1.
Teacher-guided refit	One epoch; content LR 10^{-6}; head LR 10^{-4}; $\beta = 0.1$; record replacement probability 0.25.
Verifier	64 epochs; 32 sources per batch; hidden LR 10^{-5}, readout LR 10^{-3}; auxiliary-target mixture 0.5.
Sampling	DLM temperature 0.7; diffusion start $t_\star = 800$; up to 8 proposals sharing 80 DLM calls, including verification and revision.

B Periodic Inference and Construction

Factors and evidence. Number sites from one. The fixed tree has root 1 and parent $\lfloor j/2 \rfloor$ for child $j > 1$. For H harmonics and bound $B > 0$, the periodic score is

$$g_{\psi,ij}(\delta; s) = B \tanh\left[\frac{1}{\sqrt{H}} \sum_{h=1}^{H} \left(c_{ij,h}(s) \cos \frac{2\pi h \delta}{Q} + d_{ij,h}(s) \sin \frac{2\pi h \delta}{Q} \right)\right]. \tag{B.1}$$

In Eq. (6), φ collects $H^{-1/2}$-scaled cosine and sine features for harmonics $1, \ldots, H$, and $a_{\psi,ij}$ collects the corresponding coefficients $c_{ij,h}$ and $d_{ij,h}$. The coefficients come from an MLP over normalized endpoint hidden states, species embeddings, six entries of $G_L / \operatorname{tr}(G_L)$, $\log|\det L|$, axis identity, and masked sine/cosine features of transverse differences. Explicit active-axis values are excluded from these geometric features; visible values enter as evidence and through contextual hidden states. For fixed coefficients, $g(\delta + Q; s) = g(\delta; s)$ and $|g(\delta; s)| < B$. These properties describe the learned conditional compatibility score used in the tree distribution.

Aliases and exact tree inference. For a fixed context, let $\ell^{\mathrm{tok}}_{\theta,i}(b; s)$ be the DLM logit of coordinate token b, and let $\mathrm{can}(b)$ map its periodic aliases to one grid value. Before temperature scaling, an unresolved site's processed unary score is

$$u_i(v) = \log \sum_{b:\mathrm{can}(b)=v} \exp \ell^{\mathrm{tok}}_{\theta,i}(b; s) = \ell_{\theta,i}(v; s), \qquad v \in \mathcal{S}_i(s).$$

This is the score used in Eq. (7); aliases already merged by the caller are counted once. Observed sites have singleton support and zero unary score at that value. With parent i and child j, upward

log messages are

$$
\begin{aligned}
b_j(w) &= u_j(w)/\tau + \sum_{k \in \mathrm{ch}(j)} m_{k \to j}(w), \\
m_{j \to i}(v) &= \log \sum_{w \in \mathcal{S}_j} \exp[b_j(w) + g_{\psi,ij}((v-w) \bmod Q; s)/\tau], \\
\log Z(s) &= \log \sum_{v \in \mathcal{S}_1} \exp b_1(v).
\end{aligned}
\tag{B.2}
$$

Each message sums the contributions of the entire child subtree conditional on its parent value; induction from the leaves therefore gives the exact partition function at the root. Sample the root using b_1, then each child using $b_j(w) + g_{\psi,ij}((z_i - w) \bmod Q; s)/\tau$. Conditional independence of child subtrees makes this ancestral draw exact for the fixed tree law. There are $N-1$ edges and at most Q^2 state pairs per edge, giving $O(NQ^2)$ work for one fixed tree distribution. Tree message passing computes its normalizer and marginals and supports candidate sampling. Generation recomputes the conditional factors as the visible context changes. Nonempty finite supports with finite scores on their allowed states ensure $0 < Z < \infty$. An empty support produces a recorded failure.

Effect on coordinate combinations. Fix the context s, the nonempty support Ω_s, the temperature $\tau > 0$, and the processed unary scores. Let $q_0(z \mid s) \propto \exp(\sum_i \ell_{\theta,i}(z_i; s)/\tau)$ be the distribution with the interactions removed, and write $R_\psi(z; s) = \sum_{(i,j) \in \mathcal{T}} g_{\psi,ij}(z_i - z_j; s)$ and $q_\psi = q_{\theta,\psi}$. For two assignments $z, z' \in \Omega_s$ with finite scores,

$$
\log \frac{q_\psi(z \mid s)}{q_\psi(z' \mid s)} - \log \frac{q_0(z \mid s)}{q_0(z' \mid s)} = \frac{R_\psi(z; s) - R_\psi(z'; s)}{\tau}. \tag{B.3}
$$

Substituting Eq. (7) cancels each distribution's normalizer within its probability ratio and then cancels the unary terms across ratios. If $R_\psi(z; s) > R_\psi(z'; s)$, the odds of z relative to z' increase compared with the unary-only distribution. This isolates how periodic compatibility changes combination preferences while individual-field scores remain fixed. Lattice-conditioned relations therefore enter coordinate choice during construction. The likelihood gradient below connects these preferences to the assignments observed in training crystals.

Likelihood and its gradient. For fixed s, a target $z^* \in \Omega_s$, and support independent of ψ, the per-example loss in Eq. (8) is $\log Z_{\theta,\psi}(s) - \Phi_{\theta,\psi}(z^*; s)/\tau$. Differentiating the finite sum defining Z gives

$$
\nabla_\psi[-\log q_{\theta,\psi}(z^* \mid s)] = \frac{1}{\tau}\left(\mathbb{E}_{z \sim q_{\theta,\psi}(\cdot \mid s)}[\nabla_\psi \Phi_{\theta,\psi}(z; s)] - \nabla_\psi \Phi_{\theta,\psi}(z^*; s)\right). \tag{B.4}
$$

With the constructor frozen, only the compatibility scores contribute gradients. The gradient compares the observed assignment with the model expectation, whose normalization and marginals are computed by tree inference. Observed crystal assignments thus supervise the relation scores that alter the coordinate-combination odds in Eq. (B.3).

Training views and decoding. Relation training samples an axis uniformly and a prefix length uniformly from $0, \ldots, N-1$. Count, species, lattice, earlier axes, and the current-axis prefix are visible; later axes and the remaining sites on the active axis are masked. The clean assignment on that axis is the likelihood target, and only ψ is updated. Generation completes lattice fields first, followed by coordinate axes and element blocks. The decoder uses the base DLM's confidence to order completion and updates the visible context as geometry is determined. Confidence is computed from the untempered unary probability of a candidate value. Training uses typed supports; generation additionally applies legality checks against completed visible sites, represented as unary restrictions at the current context.

C FEEDBACK TARGETS AND LEARNING OBJECTIVES

Collecting physical feedback. The reference implementation collects feedback from conditions marked for training, using periodic construction and the frozen diffusion model to generate the reference structures. During a single offline collection, references and executed candidates undergo the same CHGNet relaxation and checking protocol (Deng et al., 2023). Each record preserves the reference, proposed replacement fields, executed geometry, and evaluated outcomes. Records retain their source identifier and training-split label; reconstruction and verifier fitting accept training records, with verifier comparisons grouped by reference. The reconstruction and verifier fits reuse these associations for different supervision targets. Novelty is evaluated on the submitted structure and energetic labels on its relaxation endpoint, following Appendix E. Outcome labels determine which proposed replacement fields supervise reconstruction.

Scope generation and decision supervision. The reconstruction heads predict a mode, a local site count, and site scores from the reference view. The first candidate uses the highest-scoring non-identity mode; a local mode selects a feasible count in $\{1, 2, 4, 8\}$ and the corresponding highest-scoring sites. Further candidates use single-site scopes organized by the learned site ranking. For one-atom cells, a full-cell scope permits changes beyond rigid translation. The unchanged reference participates separately in final verifier selection.

For a supervision record b, let m_b be its mode label, c_b its local-count label, $\mathcal{I}_b$ its selected sites, and e_b its acceptance label. Writing ℓ_{mode} and ℓ_{count} for categorical cross-entropies and ℓ_{acc} for binary cross-entropy, the feedback decision objective is

$$\ell_{\mathrm{site}}(b) = -\frac{1}{|\mathcal{I}_b|} \sum_{i \in \mathcal{I}_b} \log \pi_{\mathrm{site}}(i \mid b),$$
$$\mathcal{L}_{\mathrm{dec}} = \mathbb{E}_b\big[\omega_b(\ell_{\mathrm{mode}} + \ell_{\mathrm{acc}}) + \mathbf{1}_{\{m_b=\mathrm{local}\}}(\ell_{\mathrm{count}} + \ell_{\mathrm{site}})\big]. \tag{C.1}$$

The site distribution is a softmax over valid sites, and its loss is evaluated only for a nonempty local target. Unavailable label terms contribute zero; averaging uses all records in the minibatch. The weight ω_b is two for records flagged as confirmed stable, unique, and novel, and one otherwise. These losses use detached reconstruction features and update the decision heads. The acceptance head provides auxiliary supervision; the relative verifier selects the deployed output.

Executing modifications on a continuous reference. The reconstruction model proposes replacement tokens $\hat{y}^{(k)}$ on scope $\mathcal{M}_k$ relative to $y^0 = \mathcal{Q}(C^{\mathrm{ref}})$. Let $J_k = \{j \in \mathcal{M}_k : \mathrm{can}(\hat{y}_j^{(k)}) \neq \mathrm{can}(y_j^0)\}$ identify actual changes, where can merges periodic aliases. For a numerical-field reader ξ_j and token decoder D_j, execution is

$$\xi_j(C^{(k)}) = \begin{cases} D_j(\hat{y}_j^{(k)}), & j \in J_k, \\ \xi_j(C^{\mathrm{ref}}), & j \notin J_k, \end{cases} \qquad C^{(0)} = C^{\mathrm{ref}}. \tag{C.2}$$

Thus unchanged fields retain continuous precision: a native coordinate 0.237 encoded as 0.24 remains 0.237 if the proposal predicts 0.24 again; a proposal of 0.25 replaces that field with 0.25. Equation (A.1) gives the Cartesian effect of coordinate and lattice changes, including lattice updates with preserved fractional coordinates. Scope, actual changes, proposal tokens, and executed geometry are kept as separate attributes of b_k, so physical labels are associated with the submitted modification. The unchanged reference is the identity record b_0 (KEEP).

Physical comparisons at fixed composition. Each executed modification retains the reference species and atom count. Following the composition–structure distinction in CrysLLMGen (Khastagir et al., 2025), consider successfully evaluated structures under the same energy model and fixed hull reference. Let $e_{\mathrm{rel}}(C)$ be the per-atom energy at the relaxation endpoint and $h(A)$ the hull energy at that composition, expressed on the same energy scale. Then

$$E_{\mathrm{hull}}(C) = e_{\mathrm{rel}}(C) - h(A),$$
$$E_{\mathrm{hull}}(C^{(k)}) - E_{\mathrm{hull}}(C^{\mathrm{ref}}) = e_{\mathrm{rel}}(C^{(k)}) - e_{\mathrm{rel}}(C^{\mathrm{ref}}). \tag{C.3}$$

The common composition baseline cancels on subtraction, so the energy difference compares the evaluated consequences of geometric modifications at fixed composition. Together with novelty

labels, these outcomes determine which replacement records supervise reconstruction and how they are weighted below.

Outcome labels and utility. For a resolved outcome, $t(C) \in \{0,1\}^2$ records whether the submitted structure is novel and its relaxation endpoint is stable or metastable, respectively; the metastable set includes stable structures. We use the normalized physical-quality utility

$$u(C) = w^\top t(C), \qquad w = (2/3, 1/3)^\top. \tag{C.4}$$

The teacher uses $\Delta u_k = u(C^{(k)}) - u(C^{\mathrm{ref}})$. Uniqueness is evaluated on the final output collection.

Original supervision. Write $E_k = E_{\mathrm{hull}}(C^{(k)})$ for a verified endpoint. Rank known physical outcomes by $r(C) = 4$ for stable and novel, 3 for other stable, 2 for metastable-only, and 0 otherwise. With usable reference and candidate outcomes, an improvement satisfies

$$I_k = \mathbf{1}[r_k > r_0 \text{ or } (r_k = r_0 \neq 4 \text{ and } E_0 - E_k > 0.01 \text{ eV/atom})]\,. \tag{C.5}$$

For references confirmed invalid or failed, a verified candidate of tier at least two also qualifies. Among applied improvements, choose the maximum of $(r_k, -E_k, -|\mathcal{M}_k|, -k)$. Its proposed scope and replacement tokens supply content supervision. A usable comparison without improvement supplies an empty-scope decision and negative acceptance, with zero content loss. This tier-and-energy rule supplies the original supervision, while the utility in Eq. (C.4) weights the finite-pool teacher.

Availability and teacher records. Applied candidates remain ordered records, even when two records execute to the same geometry. A teacher pool includes identity and is eligible only when its reference and every retained candidate have usable utilities and its teacher weights are nonuniform. For each eligible source, probability 0.25 replaces the original supervision with one teacher-sampled record; otherwise the original is retained. These draws are made once before fitting. Identity supplies no numerical target. A teacher-selected candidate supplies its proposed tokens even when it is not an original-rule improvement; its acceptance label still follows Eq. (C.5). Confirmed invalidity has zero utility. Unverified non-invalid outcomes are excluded from the teacher. Endpoint regression additionally permits a zero target for non-convergence with a reported hull, so target availability is checked separately for each objective.

Field reconstruction and reference regularization. The regularizer in Eq. (13) is

$$\mathcal{R}_{\mathrm{KL}}(\eta) = \mathbb{E}_{\mathcal{D}_{\mathrm{fb}}}\left[\frac{1}{m_\nu}\sum_{j\in\mathcal{M}_\nu} \mathrm{KL}(p_{\mathrm{ref},j}(\cdot \mid \nu)\|p_{\eta,j}(\cdot \mid \nu))\right]. \tag{C.6}$$

Here $p_{\mathrm{ref},j}$ is the field distribution of a reconstruction checkpoint frozen at each fitting stage; both distributions use the same field-specific support. Empty-target records give zero content loss but remain in the record-average denominator, with available labels still supervising decision heads.

C.1 Finite-Pool Optimization and Conditional Projection

Teacher optimum. For a positive prior p_0 over retained records, define

$$q_{T,p_0}(k) = \frac{p_0(k)\exp(\Delta u_k/\beta)}{Z_{T,p_0}}, \qquad Z_{T,p_0} = \sum_l p_0(l)\exp(\Delta u_l/\beta).$$

The teacher in Eq. (12) is the special case $p_0(k) = 1/(K+1)$ used by the implementation. For fixed utility labels and $\beta > 0$,

$$q_{T,p_0} = \arg\max_{q\in\Delta(\mathcal{B})} \{\mathbb{E}_q[\Delta u] - \beta\,\mathrm{KL}(q\|p_0)\}\,, \tag{C.7}$$

where $\Delta(\mathcal{B})$ is the probability simplex over records. Direct substitution gives

$$\mathbb{E}_q[\Delta u] - \beta\,\mathrm{KL}(q\|p_0) = \beta\log Z_{T,p_0} - \beta\,\mathrm{KL}(q\|q_{T,p_0}). \tag{C.8}$$

Nonnegative KL proves Eq. (C.7) on the retained finite pool. The common reference utility cancels in normalization, while including identity makes the unchanged structure an available target. Finite temperature yields a soft distribution over the evaluated alternatives, including lower-utility records. Physical utility thus determines the distribution over evaluated modifications from which supervision is drawn. The record mixture and partial-view rules below turn this preference into field-level training targets.

Conditional projection. For a record with $m > 0$ ordered target fields, training reveals a prefix of length c: $\Pr(c = 0) = 1/2 + 1/(2m)$ and $\Pr(c = l) = 1/(2m)$ for $1 \le l < m$. All remaining fields receive the average loss in Eq. (13). Let b be drawn from the empirical supervision records at the current fitting stage, including the mixture of original and teacher-selected records in the refit. Writing $\nu(b, c)$ for the resulting context and $y^\star_{b,j}$ for a field target, define the weighted empirical measure

$$\mu(\nu, j, y) = \mathbb{E}_{b,c}\left[\frac{\mathbf{1}_{\{\nu(b,c)=\nu,\ j\in\mathcal{M}_{\nu(b,c)},\ y^\star_{b,j}=y\}}}{m_{\nu(b,c)}}\right].$$

For each context/field pair with $a(\nu, j) = \sum_y \mu(\nu, j, y) > 0$, its conditional target is $d_{\mathrm{fb}}(y \mid \nu, j) = \mu(\nu, j, y)/a(\nu, j)$. This measure accounts for record sampling, random visible prefixes, and the inverse number of remaining fields; empty-target records contribute zero mass. At a fixed such pair, abbreviate $d = d_{\mathrm{fb}}(\cdot \mid \nu, j)$, $p_r = p_{\mathrm{ref},j}(\cdot \mid \nu)$, and $p = p_{\eta,j}(\cdot \mid \nu)$. Its contribution to the reconstruction objective is $a(\nu, j)[H(d, p) + \lambda \,\mathrm{KL}(p_r \| p)]$, where $H(d, p)$ is cross-entropy. For $\widetilde{d} = (d + \lambda p_r)/(1 + \lambda)$, expanding this cross-entropy and the forward KL gives

$$H(d, p) + \lambda \,\mathrm{KL}(p_r \| p) = (1 + \lambda)\,\mathrm{KL}(\widetilde{d} \| p) + (1 + \lambda)H(\widetilde{d}) - \lambda H(p_r). \tag{C.9}$$

The last two terms are constant in p, so the unconstrained conditional optimum is

$$p^\star(\cdot \mid \nu, j) = \frac{d_{\mathrm{fb}}(\cdot \mid \nu, j) + \lambda p_{\mathrm{ref},j}(\cdot \mid \nu)}{1 + \lambda}.$$

The target d_{fb} carries the effect of measured outcomes through record selection and weighting, while the regularizer retains a contribution from the frozen reconstruction checkpoint. These distributions concern replacement values for the permitted fields of X and L. Fitting them with shared DLM parameters makes experience from evaluated modifications available when predicting replacements for later reference structures. The relative supervision below uses the evaluated modifications to learn how to assess the resulting candidates.

C.2 Relative Verifier

Features and normalized readout. For record b_k, the deployed feature vector is

$$h_k = H_\eta(P, \mathcal{Q}(C^{\mathrm{ref}}), \hat{y}^{(k)}, a_k) \oplus d(C^{\mathrm{ref}}, C^{(k)}, a_k), \tag{C.10}$$

Here H_η denotes frozen reconstruction features. To specify the nine geometric features in d, let $\mathcal{M}_k$ be the numerical-field scope, $\mathcal{I}_k$ its coordinate sites, and $\Delta x_i = x_i^{(k)} - x_i^{\mathrm{ref}} - \mathrm{round}(x_i^{(k)} - x_i^{\mathrm{ref}})$ the componentwise wrapped displacement between corresponding sites. The features are $N/20$, $|\mathcal{I}_k|/N$, $|\mathcal{M}_k|/(6 + 3N)$, an indicator that the scope includes a lattice field, the root-mean-square and maximum of $\|\Delta x_i\|_2$, the root-mean-square and maximum of $\|\Delta x_i L^{\mathrm{ref}}\|_2/10$, and a geometry-availability flag. Both summaries run over all N sites, with lattice lengths expressed in Å. The flag is one when the paired geometries can be read with matching species order; otherwise the four displacement features and the flag are zero. Physical labels are stored separately for supervision. The verifier applies a learned SiLU projection to H_η, concatenates d, standardizes the result using fixed training statistics, and applies a final affine readout with weight matrix W_v. The verifier outputs real-valued outcome scores $v_\phi(h_k) \in \mathbb{R}^2$ and uses the same normalized weights as the utility:

$$f_\phi(h_k) = w^\top v_\phi(h_k), \qquad \widehat{g}_k = w^\top \Delta v_k, \qquad \Delta v_k = v_\phi(h_k) - v_\phi(h_0). \tag{C.11}$$

Selection with a reference alternative. Let $\mathcal{K}_{\mathrm{exec}}$ index executable candidates and include the identity record with index zero. The base selection is

$$k^* \in \arg\max_{k\in\{0\}\cup\mathcal{K}_{\mathrm{exec}}} \widehat{g}_k = \arg\max_{k\in\{0\}\cup\mathcal{K}_{\mathrm{exec}}} f_\phi(h_k). \tag{C.12}$$

The equality follows because $f_\phi(h_0)$ is constant across the alternatives. Including identity permits retention, and the relative score expresses this option at zero. Ties use the larger first-component gain, then identity, then input order.

Reference-relative supervision. An auxiliary outcome teacher T_χ uses measured reference properties during offline training to smooth candidate targets. Let $t_k = t(C^{(k)})$, $t_0 = t(C^{\mathrm{ref}})$, and $\bar{t}_k = (1-\alpha)t_k + \alpha T_\chi(k)$ with $\alpha = 0.5$; the reference target remains t_0. For each reference s with $n_s > 0$ usable applied candidates, the relative regression objective is

$$\mathcal{L}_{\mathrm{rel}}(\phi) = \mathbb{E}_s\left[\frac{1}{n_s}\sum_k \frac{1}{2}\left\|\Delta v_k - (\bar{t}_k - t_0)\right\|_2^2\right]. \tag{C.13}$$

Subtracting each reference target aligns supervision with the change from that starting structure; the learned Δv_k supplies the comparison used in Eq. (C.12). Writing $v_k = v_\phi(h_k)$ and $e(a,b) = \frac{1}{2}\|a-b\|_2^2$, the complete objective adds anchoring and within-reference ranking:

$$\mathcal{L}_{\mathrm{anc},s} = e(v_0, t_0) + \frac{1}{n_s}\sum_k e(v_k, \bar{t}_k),$$
$$\mathcal{L}_v = \mathcal{L}_{\mathrm{rel}} + \mathbb{E}_s[0.1\mathcal{L}_{\mathrm{anc},s} + 0.2\mathcal{R}_s] + 0.1\|W_v\|_F^2. \tag{C.14}$$

For the usable candidates of reference s, including identity, define $\mathcal{P}_s = \{(i,j) : t_{i,1} > t_{j,1}\}$. Then $\mathcal{R}_s$ is the mean of $\mathrm{softplus}(-[\Delta v_{i,1} - \Delta v_{j,1}])$ over $\mathcal{P}_s$, and is zero if this set is empty. The penalty on W_v regularizes the final readout defined above. References receive equal total weight. The auxiliary teacher uses endpoint binary cross-entropy and within-reference ranking, with its first output bounded by the second. Deployment uses the reconstruction and geometry features in Eq. (C.10).

D Training, Sampling, and Bounded Revision

D.1 Training and Sampling

Model initialization. The planner initializes from Meta-Llama-3-8B, and the construction DLM from LLaDA-8B-Instruct (AI@Meta, 2024; Nie et al., 2025). The reconstruction branch starts from a separate copy of the adapted crystal DLM and adds reference-conditioning modules and modification-decision heads. The diffusion model is trained independently on crystal geometry.

Training. Masked adaptation first updates the constructor's LoRA parameters and vocabulary tables; the periodic head is then fitted with the constructor frozen. The reconstruction copy receives geometric warm-up, an initial feedback fit, and a teacher-guided refit, with both feedback fits reusing one offline collection of evaluated modifications. During the initial feedback fit, numerical prediction updates LoRA and conditioning modules, while decision losses update their heads from detached features. The teacher-guided refit updates LoRA and decision heads with conditioning modules fixed. The verifier is fitted afterward on frozen features from the final reconstruction model. The planner, construction DLM, periodic head, and diffusion model remain fixed throughout feedback learning. Appendix A and Table 5 give the module objectives, optimization settings, and training scales; Appendix B specifies the periodic training views and inference.

Sampling. A sampled planner condition fixes A and N; the construction DLM completes L and then X, with the periodic head using A, L as conditions for coordinate combinations. Diffusion updates (X, L) to form the reference structure while preserving A. The reference protocol physically evaluates these references and retains those confirmed stable, unique, and novel. Retained references remain part of the final output collection for uniqueness evaluation. For other editable references, reconstruction modifies X within coordinate scopes and permits changes to both X, L in the full-cell scope, with A, N fixed. The verifier selects among executed candidates and the reference; at most one additional field revision uses the shared call budget. The selected geometry is saved as the generated output and subsequently relaxed for physical evaluation. All fitted models remain fixed during sampling. Appendices B and C.2 give the construction and selection rules, and Appendix D.2 specifies the optional revision gates and shared budget.

D.2 Bounded Revision

One optional field revision. After base selection, revision is allowed only for an applied single-site local-coordinate winner whose actual changes affect one site and preserve the lattice. A risk

predictor is trained on measured invalid/non-converged outcomes versus verified endpoints; verified but energetically unstable outcomes are not positive risk labels. It uses geometry/contact features at deployment. Under the shared 80-call budget, inspect at most three changed coordinates. For a remasked field j, let $p_j(z)$ be its legal conditional and r_{jz} the predicted risk of its executed counterfactual. Use

$$q_{j,\lambda_r}(z) \propto p_j(z)e^{-\lambda_r r_{jz}}, \qquad \mathrm{KL}(q_{j,\lambda_r}\|p_j) \le \kappa. \tag{D.1}$$

Reduce the requested strength if necessary to satisfy the KL budget. The default strength is 2 and $\kappa = 0.1$. The winner's risk must exceed the configured gate, and its predicted decrease relative to the best tilted field must exceed the same gate. The reference default is approximately 0.188666; its calibration is a run-level choice. Sample at most one revised value. After a passed intervention gate, rerank identity, the original winner, and any new executable candidate by $\widehat{g}_k - \frac{1}{6}[\widehat{r}_k - \widehat{r}_0]$. A failed eligibility, support, budget, or benefit check leaves the original selection unchanged. The trained models remain fixed during this bounded revision.

Risk tilting and the divergence budget. For one field, fix a nonempty finite legal support with $p(z) > 0$ and finite predicted risks $r(z)$. The normalized tilted distribution is $q_{\lambda_r}(z) = p(z)e^{-\lambda_r r(z)}/Z_r(\lambda_r)$, where $Z_r(\lambda_r) = \sum_z p(z)e^{-\lambda_r r(z)}$. For $\lambda_r > 0$ and any distribution q on this support, substitution gives

$$\mathbb{E}_q[r] + \lambda_r^{-1}\,\mathrm{KL}(q\|p) = \lambda_r^{-1}\,\mathrm{KL}(q\|q_{\lambda_r}) - \lambda_r^{-1}\log Z_r(\lambda_r). \tag{D.2}$$

Hence q_{λ_r} minimizes predicted risk penalized by divergence from the original field conditional. Also, $\frac{d}{d\lambda_r}\log Z_r = -\mathbb{E}_{q_{\lambda_r}}[r]$ and $\frac{d}{d\lambda_r}\mathbb{E}_{q_{\lambda_r}}[r] = -\operatorname{Var}_{q_{\lambda_r}}(r)$, so

$$\begin{aligned} \mathrm{KL}(q_{\lambda_r}\|p) &= -\lambda_r\mathbb{E}_{q_{\lambda_r}}[r] - \log Z_r(\lambda_r), \\ \frac{d}{d\lambda_r}\,\mathrm{KL}(q_{\lambda_r}\|p) &= \lambda_r \operatorname{Var}_{q_{\lambda_r}}(r) \ge 0. \end{aligned} \tag{D.3}$$

Because $q_0 = p$ and the divergence is continuous, reducing λ_r toward zero enforces the KL budget in Eq. (D.1).

D.3 Positioning the DLM Design

CrystalDLM adopts a pretrained masked DLM because observed-field conditioning supports both partial construction and reference-conditioned replacement (Sahoo et al., 2024; Nie et al., 2025). The crystal interface uses typed field supports, while the periodic head shapes coordinate candidates through lattice-conditioned interactions. The feedback model reuses this conditional interface with its own parameters and outcome-derived supervision. The sampling settings in Table 5 bound the number of calls used for proposals and optional revision.

E Evaluation Protocol

Evaluation cohorts and parsing. Complete-generation and draft-construction comparisons each collect $n_{\mathrm{gen}} = 10{,}000$ parseable structures per dataset and method. Failed CIF generations are discarded and sampling continues until the collection target is met. If n_{attempt} denotes all generation attempts, sampling success is $100n_{\mathrm{gen}}/n_{\mathrm{attempt}}$. Structural and compositional validity are computed over the collected cohort. The joint V.U.N., S.U.N., and M.S.U.N. evaluation uses a separate cohort of $n_{\mathrm{eval}} = 5{,}000$ structures per method, evaluated independently of the complete-generation cohort. Coverage recall is normalized by the reference-set size, and property distances compare distributions over the relevant evaluable subsets.

Validity checks. Following the benchmark conventions of Xie et al. (2022), the structural check rejects interatomic distances below 0.5 Å or cell volumes below 0.1 Å^3. Compositional validity uses charge-balance and electronegativity screening with SMACT (Davies et al., 2019), including the mixed-valence supplement in the reference implementation. Each validity percentage averages the corresponding binary check over the collected structures. Parsing success determines entry to this collected cohort; validity is checked subsequently.

Fingerprint coverage. The Coverage columns report Precision and Recall. Precision measures the fraction of generated structures near the reference set, whereas Recall measures the fraction of reference structures covered by generated structures. Higher values indicate closer reference agreement and broader reference coverage, respectively. Let $\mathcal{G}_{\mathrm{fp}}$ contain generated structures with both fingerprints available, and let $\mathcal{R}$ be the fingerprinted reference set. Structural distance d_s uses mean site-wise CrystalNN fingerprints, and composition distance d_c uses standardized Magpie descriptors (Xie et al., 2022; Ward et al., 2016). For a query C and nonempty comparison set $\mathcal{S}$, define

$$\chi(C;\mathcal{S}) = \mathbf{1}\left[\min_{D\in\mathcal{S}} d_s(C,D) \le \delta_s\right] \mathbf{1}\left[\min_{D\in\mathcal{S}} d_c(C,D) \le \delta_c\right],$$
$$\text{Precision} = \frac{100}{n_{\mathrm{gen}}} \sum_{C\in\mathcal{G}_{\mathrm{fp}}} \chi(C;\mathcal{R}), \qquad \text{Recall} = \frac{100}{|\mathcal{R}|} \sum_{C\in\mathcal{R}} \chi(C;\mathcal{G}_{\mathrm{fp}}). \tag{E.1}$$

The cutoffs δ_s and δ_c specify tolerances in structural- and compositional-fingerprint space. A query is covered when both nearest-neighbor distances satisfy their respective cutoffs; the two nearest neighbors may correspond to different crystals. Precision applies this criterion to generated structures, whereas Recall applies it to reference structures. For fixed fingerprints and sample collections, increasing either cutoff can only increase or preserve both coverage rates, so comparisons between methods on the same dataset require common cutoffs. The MP-20 reference cutoffs are $\delta_s = 0.4$ and $\delta_c = 10.0$. Missing generated fingerprints do not contribute covered events but remain in the precision denominator; empty generated support gives zero coverage.

Property-distribution distances. The reference evaluator compares properties of the fingerprint-valid generated subset with the reference distribution. Mass density ρ is mass per unit-cell volume in g/cm^3, and $N_{\mathrm{el}} = |\{a_i\}_{i=1}^{N}|$ counts distinct chemical elements, rather than total atoms N. For generated and reference empirical cumulative distributions F_G and F_R, the one-dimensional Wasserstein-1 distance is

$$W_1(F_G, F_R) = \int_{-\infty}^{\infty} |F_G(u) - F_R(u)|\, du. \tag{E.2}$$

For property $p \in \{\rho, N_{\mathrm{el}}\}$, Tables 1 and 3 report $\mathrm{wdist}(p) = W_1(F_G^p, F_R^p)$, where superscripts identify the property. The distance has that property's units. Smaller values indicate closer distributions, not lower density or fewer elements; zero denotes identical empirical distributions.

Physical screening and structure matching. The reference evaluator uses CHGNet 0.3.0 (Deng et al., 2023) with joint position/cell FIRE relaxation, at most 1,000 steps, force tolerance 0.1 eV/Å, and stress tolerance 0.5 GPa. Endpoint geometry, convergence, and finite energy are checked. The proxy hull difference compares relaxed CHGNet energy per atom with the Materials Project GGA/GGA+U competing-phase hull at the same composition (Jain et al., 2013). We use this potential-based energy difference as the screening estimate of E_{hull}.

Stability and metastability concern the relaxation endpoints, with thresholds $E_{\mathrm{hull}} \le 0$ and $E_{\mathrm{hull}} \le 0.1$ eV/atom, respectively.

Uniqueness and novelty. Both indicators are evaluated on submitted structures before relaxation. We compare periodic atomic arrangements with pymatgen's `StructureMatcher` (Ong et al., 2013), using a fractional lattice-length tolerance of 0.2, a site tolerance of 0.3 normalized by the average free length per atom, and an angular tolerance of 5°. A submitted structure is unique if it does not match any earlier output in the saved generation order. For novelty, we compare it with training-reference structures of the same reduced composition and mark it as novel if no structural match is found. The composition group identifies which structures to compare; the structural match determines whether a candidate is already represented.

Joint quality rates. V.U.N. counts structures passing both geometric and compositional validity together with uniqueness and novelty. S.U.N. and M.S.U.N. count unique, novel structures whose relaxation endpoints meet the stable and metastable thresholds, respectively. The metastable set includes the stable set; these two energetic rates use the stated endpoint screening criterion. For each corresponding event set $\mathcal{E}$, the rate is

$$\mathrm{rate}(\mathcal{E}) = 100\,|\mathcal{E}|/n_{\mathrm{eval}}, \qquad n_{\mathrm{eval}} = 5{,}000. \tag{E.3}$$

Unresolved evaluations remain in the full cohort and are separate from confirmed successes. If their count is $n_?$ and the confirmed-success count is n_+, the compatible interval is $100[n_+, n_+ + n_?]/n_{\mathrm{eval}}$.

Ablation controls. The ablations in Table 4 hold the planner, condition list, constructor checkpoint, and diffusion model fixed. DLM uses the base constructor; Periodic DLM adds periodic coordinate interactions; CrystalDLM further adds the complete feedback stage. Structural validity is measured on 10,000 outputs per variant, and S.U.N./M.S.U.N. on 5,000. The latter two variants share reference identities. Geometry is checked on saved outputs before physical relaxation, while energetic outcomes use relaxation endpoints. Uniqueness and novelty are recomputed for each variant's final collection. The feedback comparison measures the combined contribution of learned proposals, selection, and the associated inference operations.

Datasets and reporting. MP-20 and Perov-5 follow the benchmark data introduced with CDVAE (Xie et al., 2022); MPTS-52 is the larger-cell, time-split Materials Project benchmark (Baird et al., 2024). Tables 1, 2, 3, and 4 report point estimates under the collection protocols above. The draft comparison evaluates the initial language-model construction stage. Draft and complete-pipeline results describe their respective output collections, while Table 4 uses the shared-reference comparison specified above.