

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 MIAD: MIRAGE ATOM DIFFUSION FOR DE NOVO CRYSTAL GENERATION

Anonymous authors

Paper under double-blind review

ABSTRACT

In recent years, diffusion-based models have demonstrated exceptional performance in searching for simultaneously stable, unique, and novel (S.U.N.) crystalline materials. However, most of these models don't have the ability to change the number of atoms in the crystal during the generation process, which limits the variability of model sampling trajectories. In this paper, we demonstrate the severity of this restriction and introduce a simple yet powerful technique, mirage infusion, which enables diffusion models to change the state of the atoms that make up the crystal from existent to non-existent (mirage) and vice versa. We show that this technique improves model quality by up to $\times 2.5$ compared to the same model without this modification. The resulting model, Mirage Atom Diffusion (MiAD), is an equivariant joint diffusion model for de novo crystal generation that is capable of altering the number of atoms during the generation process. MiAD achieves an 8.2% S.U.N. rate on the MP-20 dataset, which substantially exceeds existing state-of-the-art approaches.

1 INTRODUCTION

Deep generative modeling opens new horizons of possibilities across various fields. Its impact is particularly profound in the natural sciences, where generative models can act as powerful exploration engines (Jumper et al., 2021) capable of dramatically accelerating scientific and technological progress. One area where this potential is especially evident is materials science. Despite the vast number of hypothetically possible materials, only a small fraction are capable of stable existence under physical constraints, the underlying principles of which remain largely unknown; this poses a significant challenge to the discovery of new materials.

Several approaches to deep generative modeling, including Generative Adversarial Networks (GANs) (Goodfellow et al., 2014; Karras et al., 2019), Variational Autoencoders (VAEs) (Kingma & Welling, 2013), Large Language Models (LLMs) (Devlin et al., 2019; Touvron et al., 2023), as well as Diffusion models (Sohl-Dickstein et al., 2015; Ho et al., 2020; Song et al., 2020; Vahdat et al., 2021) and Flow-Matching models (Lipman et al., 2022) shown strong performance in generating complex objects. The latter two paradigms have demonstrated their potential in handling the trade-off between object quality and generation diversity (Xiao et al., 2022). These models are the same in organizing the generation process as a trainable, iterative process, which progressively refines the output over a sequence of steps.

Applying diffusion models to data from the natural sciences is non-trivial due to the need to encapsulate specific properties within a trainable dynamic framework. Multiple approaches are tailored for this objective, each addressing a particular class of properties. Currently, diffusion models can operate in non-Euclidean spaces (Huang et al., 2022; De Bortoli et al., 2022; Chen & Lipman, 2023; Okhotin et al., 2023; Hoogeboom et al., 2021) and exhibit symmetrical properties (Hoogeboom et al., 2022; Klein et al., 2023).

In recent years, diffusion models have demonstrated their potential in the domain of material science. The DiffCSP model (Jiao et al., 2024a) has shown ability to generate 3D crystal structures while effectively capturing the intrinsic regularities of the crystal manifold. Additionally, the MatterGen model (Zeni et al., 2024) demonstrated that diffusion models can generate stable and novel materials – an essential quality characteristic in this task (Kazeev et al., 2025). While it is natural to expect generative models to generalize beyond the training data and create novel samples, this objective

can sometimes conflict with the training loss, which often emphasizes fidelity to the observed data distribution. Such tension can negatively influence the model’s ability to explore new regions of the data space.

However, by design, these models are restricted to generating crystals with a fixed number of atoms, which limits their ability to perform intuitive operations such as adding or removing specific atoms during generation. As we show in Section 6, this constraint reduces the model’s flexibility and hinders its ability to explore a broader range of plausible crystal structures, impacting both the diversity and quality of the generated outputs.

In this paper, we present the next steps toward improving diffusion models for crystal generation and demonstrate their significance in the search for novel and stable materials:

- We propose a simple yet powerful technique, *mirage infusion*, which broadens the original space of 3D crystal structures, enabling the diffusion model to modify the number of atoms in the crystal during the generation process.
- We examine the sensitive parameters of the proposed technique and their impact on the quality of the generative model through a series of experiments.
- We demonstrate that the proposed approach significantly enhances the performance of the base joint diffusion model and substantially surpasses the previous state-of-the-art model.

2 PRELIMINARIES

2.1 DIFFUSION MODELS

A diffusion model is a generative model that consists of forward and backward processes:

$$q(x_{0:T}) = q(x_0) \prod_{t=1}^T q(x_t|x_{t-1}), \quad p_\theta(x_{0:T}) = p(x_T) \prod_{t=1}^T p_\theta(x_{t-1}|x_t),$$

where $q(x_0)$ represents the data distribution, $q(x_t|x_{t-1})$ denotes the transition kernel which gradually adds noise to data objects, $p(x_T)$ is the prior distribution from which the diffusion model begins generation, and $p_\theta(x_{t-1}|x_t)$ is the trainable transition kernel used for step-by-step object denoising during the sampling procedure. The original training objective for the backward process is the Evidence Lower Bound (ELBO). Alternatively, in continuous spaces, diffusion models can be reformulated using *score-matching*. In this framework, the model is trained to approximate the *conditional score* $\nabla_{x_t} \log q(x_t|x_0)$, while sampling relies on the *unconditional score* $\nabla_{x_t} \log q(x_t)$.

Finally, for data with complex internal structures, one can factorize objects into several components and define a separate diffusion process for each of them. As an example, crystals can be factorized into three components: lattice matrix, fractional coordinates of atoms and atom types, where each component belongs to a space with a specific geometry that differs from the others. In the model, which we extend in this work, the crystal structure is decomposed into a lattice matrix, fractional atomic coordinates, and atom types — each residing in a space with distinct geometric properties. In this setting, the sampling procedure requires simultaneous denoising of all these components. Therefore, we can use a single neural network, which takes the noisy versions of all object components and denoises them according to their respective diffusion processes. In this way, the diffusion model for each component is conditioned on the noisy versions of all object components forming a *joint diffusion model*. The resulting training objective is a weighted sum of the objectives from all diffusion processes.

2.2 CRYSTAL REPRESENTATION

The representation of a 3D crystal can be reduced to the representation of a *unit cell*, with the entire crystal being an infinite repetition of its unit cell. A unit cell is represented by the lattice $L = [l_1, l_2, l_3] \in \mathbb{R}^{3 \times 3}$ — three basis vectors that define its geometry, $F \in [0, 1]^{N_{\text{atoms}} \times 3}$ — the fractional coordinates of atoms within the unit cell, and $A \in \{1, \dots, N_{\text{types}}\}^{N_{\text{atoms}}}$ — the types of atoms. The number of atoms in the unit cell, $N_{\text{atoms}} \in \mathbb{N}_+$, varies depending on the specific crystal.

108 The entire crystal is an infinite periodic structure, defined as a set of atoms given by:
 109

$$110 \quad \{(a_i, x_i) \mid x_i = (f_i + k)L, \forall k \in \mathbb{Z}^3\},$$

112 where each atom is represented by a pair consisting of its type a_i and its Cartesian coordinates x_i .
 113 Thus, the task of generating a crystal reduces to generating the triplet $\mathcal{M} = (L, F, A)$.
 114

115 2.3 DIFFUSION MODEL FOR CRYSTALS

116 The joint diffusion model over crystal structures $\mathcal{M} = (L, F, A)$ is proposed in Jiao et al. (2024a);
 117 Zeni et al. (2024). In this model, the forward process is factorized across the lattice L , fractional
 118 coordinates F , and atoms types A :
 119

$$120 \quad q(\mathcal{M}_t | \mathcal{M}_{t-1}) = q(L_t | L_{t-1})q(F_t | F_{t-1})q(A_t | A_{t-1}),$$

121 where $\mathcal{M}_0 = (L_0, F_0, A_0)$ denotes a clean crystal, while $\mathcal{M}_t = (L_t, F_t, A_t)$ represents an interme-
 122 diate noisy version of the crystal.
 123

124 2.3.1 DIFFUSION COMPONENTS

125 **Lattice** Given that $L_0 \in \mathbb{R}^{3 \times 3}$, we can employ DDPM (Ho et al., 2020) with Gaussian distributions:
 126

$$127 \quad q(L_t | L_{t-1}) = \mathcal{N}\left(L_t | \sqrt{\beta_t}L_{t-1}, \sqrt{1 - \beta_t}I\right)$$

$$128 \quad p(L_T) = \mathcal{N}(L_T | 0, I)$$

$$129 \quad p_\theta(L_{t-1} | \mathcal{M}_t) = \mathcal{N}(L_{t-1} | \mu_\theta(\mathcal{M}_t, t), \sigma_t^2 I),$$

130 where $\beta_t \in \mathbb{R}_+$ – sets the scale of noise for step t and $p_\theta(L_{t-1} | \mathcal{M}_t)$ is conditioned on the noisy
 131 versions of all crystal components, as it is a joint diffusion model. The training objective for this
 132 diffusion component is defined as a weighted ELBO, which reduces to the following form:
 133

$$134 \quad \mathcal{L}_L = \sum_{t=2}^T \gamma_t \mathbb{E}_{\mathcal{M}_0 \sim q(\mathcal{M}_0), \mathcal{M}_t \sim q(\mathcal{M}_t | \mathcal{M}_0)} D_{\text{KL}} [q(L_{t-1} | L_t, L_0) \parallel p_\theta(L_{t-1} | \mathcal{M}_t)]$$

135 **Fractional coordinates** The space of fractional coordinates is periodic. Therefore, we need to
 136 utilize a diffusion model that can accommodate data of this nature. One possible option is to use a
 137 Wrapped Normal distribution for each coordinate of each atom:
 138

$$139 \quad q(F_t | F_{t-1}) = \mathcal{WN}(F_t | F_{t-1}, (\sigma_t^2 - \sigma_{t-1}^2)I)$$

$$140 \quad p(F_T) = \mathcal{U}(F_T | 0, 1),$$

141 where $\sigma_t \in \mathbb{R}_+$ – set the scale of noise for step t . A particularly suitable and efficient method for
 142 training the backward process in this model is to adopt Riemannian score matching (De Bortoli et al.,
 143 2022) with the following objective:
 144

$$145 \quad \mathcal{L}_F = \mathbb{E}_{\mathcal{M}_0 \sim q(\mathcal{M}_0), t \sim \mathcal{U}(1, T), \mathcal{M}_t \sim q(\mathcal{M}_t | \mathcal{M}_0)} \|\nabla_{F_t} \log q(F_t | F_0) - s_\theta(\mathcal{M}_t, t)\|_2^2,$$

146 where s_θ is a neural network that approximates the unconditional score for the generating procedure
 147 proposed by Jiao et al. (2024a)
 148

149 **Atom types** The atom type is a discrete variable drawn from a fixed set of N_{types} possible elements.
 150 We use D3PM (Austin et al., 2023) for this component:
 151

$$152 \quad q(A_{t,i} | A_{t-1,i}) = \text{Cat}(A_{t,i} | Q_t A_{t-1,i}^{\text{onehot}})$$

$$153 \quad p(A_{T,i}) = \text{Cat}(A_{T,i} | \mathbf{1} / N_{\text{types}})$$

$$154 \quad p_\theta(A_{t-1,i} | \mathcal{M}_t) = \text{Cat}(A_{t-1,i} | c_{\theta,i}(\mathcal{M}_t, t)),$$

155 where $A_{t,i}$ – the atom i in the crystal \mathcal{M}_t , $A_{t-1,i}^{\text{onehot}}$ – zero vector of size N_{types} with 1 on position
 156 $A_{t-1,i}$, $\mathbf{1}$ – unit vector of size N_{types} , $Q_t \in \mathbb{R}_+^{N_{\text{types}} \times N_{\text{types}}}$ – matrix that gradually spreads probability
 157

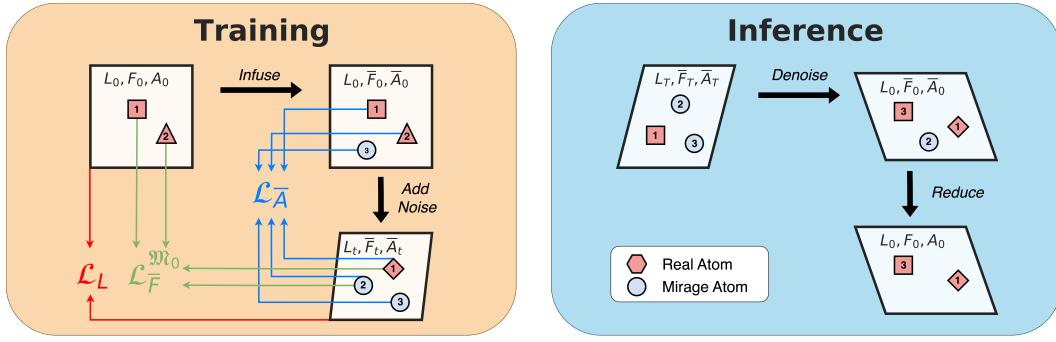


Figure 1: Overview of the proposed mirage infusion technique.

mass across all types on each step t , $c_{\theta,i}$ – prediction of probabilities of each type for atom i . The training objective takes the following form:

$$\mathcal{L}_A = \sum_{t=2}^T \mathbb{E}_{\mathcal{M}_0 \sim q(\mathcal{M}_0), \mathcal{M}_t \sim q(\mathcal{M}_t | \mathcal{M}_0)} \sum_{i=1}^{N_{\text{atoms}}} D_{\text{KL}} [q(A_{t-1,i} | A_{t,i}, A_{0,i}) || p_{\theta}(A_{t-1,i} | \mathcal{M}_t)]$$

Although the model is capable of generating crystals with different numbers of atoms, this number must be specified before the generation process begins.

The final objective for this joint diffusion model is a weighted sum of the objectives for the components L , F , and A :

$$\mathcal{L} = \kappa_1 \mathcal{L}_L + \kappa_2 \mathcal{L}_F + \kappa_3 \mathcal{L}_A \rightarrow \min_{\theta},$$

where the coefficients $\kappa_i > 0$ for $i = 1, 2, 3$ affect component prioritization, provided that we use the same neural network for making predictions.

2.3.2 INVARIANCES

The joint diffusion model outlined in the previous section is applied within the domain of crystals, which possesses distinct properties referred to as *symmetries*. To ensure these symmetries are preserved, the model must be parametrized in a particular manner. Further details regarding the specific types of symmetries present in the domain of crystals, as well as the organization of diffusion components required to achieve invariance to these symmetries, are provided in Appendix A.

3 MIRAGE INFUSION

Motivation In the joint diffusion model, the number of atoms in a crystal N_{atoms} must be fixed in advance during the generation procedure to sample \mathcal{M}_T from the prior distribution:

$$\mathcal{M}_T \sim p(\mathcal{M}_T | N_{\text{atoms}}) = p(L_T, F_T, A_T | N_{\text{atoms}}) = p(L_T)p(F_T | N_{\text{atoms}})p(A_T | N_{\text{atoms}}), \quad (1)$$

where N_{atoms} is typically drawn from a categorical distribution $p(N_{\text{atoms}})$, estimated from the training data (Jiao et al., 2024a), and it usually spans a limited range of discrete values (e.g., 1 to 20). Consequently, the number of atoms in a crystal’s unit cell is fixed during the whole process of generation, which, as we show in the following sections, crucially limits the model’s flexibility. To address this limitation, we introduce a framework that supports adding or removing atoms during generation.

Core idea Traditional diffusion models cannot change object size during generation. To enable this, we reinterpret the addition and removal of atoms as transitions between different types of atoms. Specifically, we introduce *mirage atoms* – placeholder atoms that may either materialize into real atoms or vanish during the generation process.

216 In a crystal representation \mathcal{M} , mirage atoms are distinguished by assigning them a special type.
 217 Since F is continuous, it cannot be used to flag mirage atoms directly. However, A is discrete, so we
 218 designate a new atom type 0 to represent a mirage atom.
 219

220 **Method** To support mirage atoms, we define an *expanded crystal domain* $\overline{\mathcal{M}} = (L, \overline{F}, \overline{A})$, where
 221 all objects have a fixed number of atoms N_m that is greater than or equal to the maximum number of
 222 atoms in a crystal in the training dataset. Mirage atoms are simply atoms with type 0 in this domain.
 223

224 We define two mappings:

- 225 • **Infusion:** Adds mirage atoms to a real crystal \mathcal{M} . The original atoms are kept unchanged,
 226 while the additional atoms are initialized with type 0, and fractional coordinates are drawn
 227 from the uniform distribution over the unit cell $\mathcal{U}(0, 1)^3$.
- 228 • **Reduction:** Removes mirage atoms from expanded crystal $\overline{\mathcal{M}}$ by filtering out all atoms
 229 with type 0, restoring the original representation \mathcal{M} .
 230

231 Using this setup, we can train a diffusion model in the ex-
 232 panded domain. The model architecture remains nearly
 233 the same, with the only change being an extra atom type
 234 (type 0) in the D3PM diffusion component for atom
 235 types.

236 In training, crystal from the training data \mathcal{M}_0 is infused
 237 with the mirage atoms with randomly initialized frac-
 238 tional coordinates, effectively augmenting the dataset.
 239 The lattice loss remains unchanged. The atom-type loss
 240 remains the same as well, training the model to predict
 241 real and mirage atom types in the final structure. For
 242 fractional coordinates loss, we ignore the mirage atoms
 243 defined by the mirage’s mask $\mathfrak{M}_0 = \{i \mid \overline{A}_{0,i} \neq 0, i =\overline{1, N_m}\}$ for the original crystal \mathcal{M}_0 , since they have no
 244 ground-truth positions. This allows the model to freely
 245 adjust their positions and focus learning only on real
 246 atoms. The loss is masked accordingly:
 247

$$\mathcal{L}_F^{\mathfrak{M}_0} = \mathbb{E}_{\overline{\mathcal{M}}_0 \sim q(\overline{\mathcal{M}}_0), t \sim \mathcal{U}(1, T), \overline{\mathcal{M}}_t \sim q(\overline{\mathcal{M}}_t | \overline{\mathcal{M}}_0)} \sum_{i \in \mathfrak{M}_0} \|\nabla_{\overline{F}_{t,i}} \log q(\overline{F}_t | \overline{F}_0) - s_{\theta,i}(\overline{\mathcal{M}}_t, t)\|_2^2, \quad (2)$$

250 where $q(\overline{\mathcal{M}}_0)$ is a distribution of crystals after the infusion of mirage atoms, and $s_{\theta,i}(\overline{\mathcal{M}}_t, t)$ is a
 251 score for the fractional coordinates of atom i .
 252

253 The sampling algorithm remains the same, except for two differences: (1) at the start of generation,
 254 all crystals are sampled with the same number of atoms N_m (not with $N_{\text{atoms}} \sim p(N_{\text{atoms}})$ as in 1),
 255 (2) at the end of generation we need to apply the reduction operator to project the crystal onto the
 256 original domain.

257 Importantly, this method preserves existing symmetries because lattice representations remain un-
 258 changed and mirage atoms follow the same spatial rules as real atoms.

259 We call the proposed technique *mirage infusion*. It expands the space of generative trajectories
 260 available to the model, increasing its flexibility and expressiveness. We outline the training procedure
 261 in Algorithms 1, and the sampling procedure in Algorithm 2, as well as their schematic visualization
 262 in Figure 1.
 263

264 **Discussion** There are several design choices when defining the expanded domain. In our setup, we:
 265 (1) initialize mirage atoms with uniformly random coordinates. This choice ensures that mirage atoms
 266 can occupy arbitrary positions in the unit cell and introduces variability into the generation process.
 267 (2) Mask the loss for mirage atoms during training to avoid learning from noise. This allows the
 268 model to focus on denoising atoms that exist in the final structure while giving it the freedom to learn
 269 when and how mirage atoms should transform into real ones during generation. Skipping the mask
 would force the model to predict a zero score for mirage atoms, potentially reducing expressiveness.

Algorithm 1 MiAD Training

- 1: **repeat**
- 2: Sample $t \sim \mathcal{U}(1, T)$
- 3: Sample $\mathcal{M}_0 \sim q(\mathcal{M}_0)$
- 4: Infuse mirage atoms $\mathcal{M}_0 \rightarrow \overline{\mathcal{M}}_0$
- 5: Add noise $\overline{\mathcal{M}}_t \sim q(\overline{\mathcal{M}}_t | \overline{\mathcal{M}}_0)$
- 6: Minimize $\kappa_1 \mathcal{L}_L + \kappa_2 \mathcal{L}_F^{\mathfrak{M}_0} + \kappa_3 \mathcal{L}_A$
- 7: **until** Convergence

Algorithm 2 MiAD Sampling

- 1: Sample $\overline{\mathcal{M}}_T \sim q(\overline{\mathcal{M}}_T)$
- 2: **for** $t \leftarrow T$ **to** 1 **do**
- 3: Denoise $\overline{\mathcal{M}}_{t-1} = q(\overline{\mathcal{M}}_{t-1} | \overline{\mathcal{M}}_t)$
- 4: **end for**
- 5: Reduce $\overline{\mathcal{M}}_0 \rightarrow \mathcal{M}_0$

270 A similar idea was proposed independently by Schneuing et al. (2025) for structure-based drug
 271 design. In their method, mirage atom coordinates are initialized at the object’s center of mass, and
 272 no loss masking is applied. While both methods share conceptual goals, we argue that our design is
 273 more principled and flexible. Empirical comparisons support this claim, showing that our approach
 274 achieves higher generative quality (see Appendix B.3).

276 4 RELATED WORK

278 The field of deep generative modeling in material science has been rapidly growing in recent years.
 279 State-of-the-art approaches use a multi-step procedure for generating crystals. In this review, we
 280 organize the works by their treatment of the number of atoms during this generation process.

282 Models with a constant number of atoms directly sampled from the prior distribution; their types
 283 and coordinates are gradually refined during the generation process along with the lattice. DiffCSP
 284 (Jiao et al., 2024a) is the first pure diffusion model for crystal generation, our work directly extends
 285 their approach. Most parts of this model are described in Section 2; briefly, it is a joint diffusion
 286 model that operates with the crystal space represented as a triplet of lattice, fractional coordinates,
 287 and atom types. The following models fall within the same paradigm: MatterGen (Zeni et al., 2024),
 288 FlowMM (Miller et al., 2024), CrysBFN (Wu et al., 2025), TGDMat (Das et al., 2025), CrystalFlow
 289 (Luo et al., 2025). DiffCSP++ (Jiao et al., 2024b) is a development of DiffCSP that constrains the
 290 diffusion process so that the crystal has predefined Wyckoff symmetries, for de novo generation they
 291 are sampled from the training dataset. ADiT (Joshi et al., 2025), uses a Transformer instead of a
 292 GNN as the denoising model.

294 Models with two stages, where the first stage is used to generate an intermediate crystal representation
 295 that includes the number of atoms, from which the structure is reconstructed during the second stage,
 296 where the number of atoms is fixed: CDVAE (Xie et al., 2022), FlowLLM (Sriram et al., 2024),
 297 WyFormer (Kazeev et al., 2025). All these models combine non-diffusion generative models in the
 298 first stage with diffusion-based refinement in the second stage.

299 Models that change the number of atoms during the entire generation process, i.e., the number of
 300 atoms is generated jointly with all other components of the crystal.

302 The first subgroup consists of methods for changing the representation of 3D crystal structures,
 303 which allows them to vary the number of atoms during the generation process, Uni-3DAR (Lu et al.,
 304 2025), UniMat (Yang et al., 2024), WyckoffDiff (Kelvinius et al., 2025). Autoregressive models
 305 naturally fall into this subgroup: CrystalFormer (Cao et al., 2024), CrystaLLM (Antunes et al., 2024),
 306 LLaMA-2 (Gruver et al., 2024).

308 The second subgroup consists of Crystal-GFN (AI4Science et al., 2023) and SHAFT (Nguyen et al.,
 309 2024) – policies, which are trained using reward functions and operate in a space of 3D crystal
 310 representations along with their symmetry groups.

312 The third subgroup consists of models, which indirectly change the number of atoms during generation
 313 via changing Wyckoff positions site symmetry: SymmCD (Levy et al., 2025) and SymmBFN (Ruple
 314 et al., 2025). The diffusion/flow process, however, only deals with an asymmetric unit of fixed size.

316 The proposed categorization underscores the inherent limitations of diffusion-like methods in ac-
 317 commodating the insertion and removal of atoms during the generation process. This finding is
 318 particularly noteworthy, as leading approaches such as DiffCSP, FlowLLM, WyFormer, and ADiT
 319 are either partially or entirely based on the diffusion paradigm.

320 5 METRICS

322 De novo crystal generation serves as the first step of the material discovery pipeline. Xie et al. (2022)
 323 proposed several metrics for evaluating generative models in this task: Structural and Compositional
 324 Validity, Coverage-Recall, Coverage-Precision, and Wasserstein distances between the distributions
 325 of various properties in the training set and those of crystals produced by the model. Although
 326 these metrics have become widespread tools for model comparison, they have severe limitations
 327 that constrain their applicability to analyzing the performance of modern generative models. We
 328 report MiAD’s results under these metrics, as well as our arguments against their use in their current

324 form, in Appendix E. Among existing approaches for comparing generative models in de novo
 325 crystal generation, we focus on S.U.N. (Zeni et al., 2024) — the proportion of materials that are
 326 simultaneously *stable*, *unique* and *novel*. This is one of the most reasonable measures of a generative
 327 model’s performance, as it directly quantifies the model’s utility for materials discovery.
 328

329 **Novelty** Represents the fraction of the generated materials that are not present in the training dataset.
 330 Following all the baselines, we use `StructureMatcher` from the `pymatgen` package (Ong et al.,
 331 2013) with the default parameters.
 332

333 **Uniqueness** As the number of generated materials grows, a model starts to repeat itself. Uniqueness
 334 is the fraction of unique materials among generated, also estimated with `StructureMatcher`.
 335

336 **Stability** To be useful, a material must actually exist under normal conditions. Ab initio prediction
 337 of experimental stability is an open research question (Tolborg et al., 2022; Sun et al., 2016) with
 338 various trade-offs between the accuracy and computational cost possible. Again, we follow state-of-
 339 the-art ML baselines (Miller et al., 2024; Kazeev et al., 2025; Zeni et al., 2024; Joshi et al., 2025) and
 340 use *energy above convex hull* E^{hull} as the stability measure. There are two important nuances.
 341

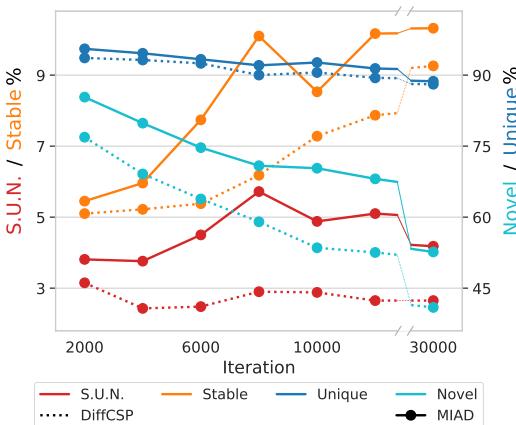
342 (1) **Stability condition** If a material has positive E^{hull} , this indicates that the same set of atoms can
 343 be rearranged into a known different configuration with a lower potential energy. This, however,
 344 does not necessarily mean that the higher-energy material will not exist under normal conditions; for
 345 example, both graphite and diamond do. We, therefore, measure the number of both *stable* materials
 346 with $E^{\text{hull}} < 0$ eV that are highly likely to exist; and *metastable* materials with $E^{\text{hull}} < 0.1$ eV, which
 347 are likely, but not certain to exist; the choice of the threshold follows Joshi et al. (2025).
 348

349 (2) **Energy computation method** Density Functional Theory (DFT) (Kohn & Sham, 1965) is the ab
 350 initio method that has been used to obtain the energy values and structures in Materials Project (Jain
 351 et al., 2013b), used both as training data and convex hull. We use DFT to support the main claims
 352 of our paper, to evaluate the performance of MiAD, and compare it to the baselines in Table 1; the
 353 computation details are described in Appendix D. Due to the large computational cost of DFT, not all
 354 baseline works use it for stability evaluation, some opt for Machine Learning Interatomic Potentials
 355 (MLIPs). MLIPs are a rapidly advancing research area, useful for stability estimation (Riebesell et al.,
 356 2023), but still fall short of DFT accuracy (Deng et al., 2025), as evident by S.U.N. values difference
 357 for the same models in Tables 1 and 2. We use two MLIP models, CHGNet (Deng et al., 2023) and
 358 eq-V2 (Barroso-Luque et al., 2024), for ablation studies and for a supplementary comparison with
 359 baseline methods for which DFT data are not available. In the cases of DFT or CHGNet, we first
 360 prerelax crystals via CHGNet for 1500 steps, while in the cases of eq-V2, we prerelax via eq-V2 for
 361 100 steps.
 362

363 6 EXPERIMENTS

364 In our experimental evaluation, we aim to highlight the importance of the model’s capability to
 365 modify the number of atoms during the generative process. The proposed approach is a combi-
 366 nation of the DiffCSP model Jiao et al. (2024a) with the proposed mirage infusion technique (see
 367 Section 3). To isolate the effect of this technique, we retain an identical neural network architecture,
 368 thereby minimizing confounding variables.
 369

370 However, it should be noted that the proposed
 371 model incurs higher computational costs, at-
 372 tributable to the increased average number of
 373 atoms per generated crystal. Additionally, mi-
 374 rage infusion necessitates a greater number of
 375 optimization iterations to attain optimal per-
 376 formance. All experiments are conducted using the
 377 MP-20 dataset (Jain et al., 2013a), and results are



378 Figure 2: Comparison of MiAD (DiffCSP with
 379 mirage infusion) and DiffCSP in terms of stabil-
 380 ity, uniqueness, novelty, and S.U.N. Stability is
 381 estimated via eq-V2.
 382

378
379 **Table 1: Crystal generation comparison via S.U.N. based on DFT.** We report metastability,
380 M.S.U.N., stability, and S.U.N. for 10 000 sampled crystals. For clarity in evaluating the model’s
381 quality, we also report the Unique&Novel rate separately for metastable and stable crystals, respec-
382 tively. MiAD outperforms all existing approaches in terms of M.S.U.N. and S.U.N. DiffCSP &
383 FlowMM results are taken from Miller et al. (2024), FlowLLM Sriram et al. (2024), ADiT Joshi et al.
384 (2025); WyFormer computed by us from the DFT structures provided by the authors.

Model	Metastability ($E^{\text{hull}} < 0.1$)			Stability ($E^{\text{hull}} < 0.0$)		
	Metastable (%) ↑	Unique&Novel (%) ↑	M.S.U.N. (%) ↑	Stable (%) ↑	Unique&Novel (%) ↑	S.U.N. (%) ↑
DiffCSP	-	-	-	5.0	66.0	3.3
FlowMM	30.6	73.5	22.5	4.6	60.9	2.8
FlowLLM	66.9	39.3	26.3	13.9	33.8	4.7
WyFormer	30.5	89.8	27.4	5.2	92.3	4.8
MP20-only ADiT	81.6	31.8	25.9	14.1	33.3	4.7
MP20-only ADiT (32M)	71.1	53.6	38.1	12.8	50.8	6.5
Jointly trained ADiT	81.0	34.8	28.2	15.4	34.4	5.3
MiAD	73.5	59.4	43.6	12.5	65.2	8.2

393
394 **Table 2: Crystal generation comparison via S.U.N. based on MLIPs.** We report stability and
395 S.U.N. estimated using CHGNet and eq-V2 for 10 000 sampled crystals. For clarity in evaluating the
396 model’s quality, we also report the Unique&Novel rate among stable crystals. MiAD outperforms all
397 existing approaches in terms of both variants of S.U.N. Results for eq-V2 are computed by us from
398 the structures provided by the authors, whereas results for CHGNet are taken from Levy et al. (2025)
399 and Ruple et al. (2025).

Model	eq-V2 ($E^{\text{hull}} < 0.0$)			CHGNet ($E^{\text{hull}} < 0.0$)		
	Stable (%) ↑	Unique&Novel (%) ↑	S.U.N. (%) ↑	Stable (%) ↑	Unique&Novel (%) ↑	S.U.N. (%) ↑
CDVAE	-	-	-	4.4	96.4	4.3
DiffCSP	3.8	66.6	2.5	11.3	78.8	8.9
DiffCSP++	2.9	73.5	2.1	11.4	75.9	8.6
FlowMM	2.0	70.4	1.4	9.1	71.7	6.5
SymmCD	2.8	70.3	2.0	9.3	73.5	6.9
SymmBFN	-	-	-	11.8	75.4	8.9
MatterGen MP-20	2.5	73.6	1.8	9.7	83.5	8.1
MiAD	9.7	57.0	5.5	19.8	65.2	12.9

409 assessed using various versions of the S.U.N. metric (see Section 5). Other experimental details are
410 provided in Appendix C.

411 Appendix B provides a detailed exposition of key design choices underlying the mirage infusion
412 technique. Specifically, we illustrate that, within the proposed definition of the expanded domain,
413 there exists flexibility in selecting N_m , the total number of real and mirage atoms for crystals in the
414 expanded domain. Moreover, mirage infusion influences the scaling of loss terms associated with
415 fractional coordinates and atom types. Our analysis reveals that suboptimal weighting of these loss
416 components in the joint diffusion model can have a nontrivial impact on overall model quality. As
417 discussed in Section 4, alternative formulations of mirage infusion have been proposed. We present
418 a comparison between our definition and a contemporary variant, underscoring the critical role of
419 expanded domain design and loss function modifications. Lastly, we compare the final version of
420 mirage infusion with the baseline model without these modifications, thereby demonstrating the
421 substantial benefits of the proposed approach simultaneously in stability, uniqueness, and novelty
422 rate (see Figure 3). This ablation is performed using S.U.N., where stability is estimated using
423 eq-V2 (Barroso-Luque et al., 2024) due to computational constraints.

424 We designate the finalized model as MiAD (Mirage Atom Diffusion) and benchmark its performance
425 against state-of-the-art methods for de novo crystal generation. Multiple versions of the S.U.N. metric
426 are employed to facilitate comprehensive comparisons across a broad range of existing models. It
427 should be noted that we do not compare against studies that lack S.U.N. metric evaluations.

428 Comparison in Table 1 presents the results, demonstrating MiAD performance against existing
429 approaches, with respect to the S.U.N. metric, where stability is assessed via DFT (Kohn & Sham,
430 1965) according to the protocol described in Appendix D. As outlined in Section 5, this configuration
431 represents the most rigorous S.U.N. evaluation and currently prevails over all other approaches. While
MiAD exhibits a lower fraction of stable crystals compared to ADiT (Joshi et al., 2025), this is

432 attributable to the tendency of this model to replicate training set samples. Consequently, stability
 433 alone is insufficient as a measure of generative model quality. At the same time, MiAD exhibits a
 434 lower fraction of unique and novel crystals among stable compared to WyFormer, while exceeding
 435 the last one in terms of stability rate. Therefore, MiAD achieves superior overall S.U.N. performance
 436 relative to all baselines (+25% relative to the closest competing method) due to the best trade-off
 437 between stability, uniqueness, and novelty. Another significant observation is that the incorporation
 438 of mirage infusion enhances the performance of the original DiffCSP model by up to $\times 2.5$ times,
 439 representing a substantial improvement.

440 Furthermore, some works only report S.U.N. computed with MLIPs, presumably due to computational
 441 constraints. MiAD comparison to them is presented in Table 2; that MiAD shows the best performance
 442 there as well.

443 Finally, we provide additional analyses of the diversity of MiAD’s generations. In Appendix F, we
 444 demonstrate that MiAD changes the number of atoms during generation and successfully produces
 445 S.U.N. crystals with varying numbers of atoms. In Appendix G, we provide the distributions of
 446 space groups for modern generative models and verify that MiAD preserves diversity in these terms.
 447 The scalability of the proposed model on larger datasets is demonstrated in Appendix H, where we
 448 compare MiAD with MatterGen on the Alex-MP20 dataset (Zeni et al., 2024).

449

450 7 DISCUSSION

451

452 Conclusion The categorization provided in Section 4 offers valuable insights into the potential
 453 applicability of the proposed technique across a range of models within the domain of de novo
 454 material generation. Beyond DiffCSP, which serves as the foundation for the model introduced
 455 in this work, the mirage infusion technique can be directly implemented, without modification,
 456 in MatterGen, FlowMM, and CrystalFlow, given that these models employ an identical crystal
 457 representation and share the same factorization of the loss function. Additionally, we hypothesize that
 458 an adapted version of this technique could be extended to models such as ADiT, CDVAE, FlowLLM,
 459 SymmCD, and SymmBFN, although further investigation is required to assess its effectiveness, and
 460 the eventual impact in these cases remains uncertain. At the same time, in Section 6 we present MiAD,
 461 demonstrating that mirage infusion gives a substantial boost in quality for the base joint diffusion
 462 model and outperforms existing state-of-the-art approaches. This improvement demonstrates the
 463 potential of approaches directed at the modification of the space with which a particular diffusion
 464 model works.

464

465

466 Limitations The proposed technique, mirage infusion, demonstrates a substantial improvement in
 467 the specific joint diffusion model for crystal generation, as evidenced by the S.U.N. metric. Neverthe-
 468 less, further research is essential to investigate various adaptations of this technique, particularly its
 469 application to other established generative models in the domain of de novo crystal generation, as
 470 well as in other areas of generative modeling, to more comprehensively assess the methodology’s
 471 position within the broader field. Additionally, the field of de novo crystal generation requires the
 472 development of new metrics capable of evaluating the diversity of existing generative models from a
 473 wider range of perspectives, thereby enhancing the understanding of the properties of both existing
 474 and newly generative models. At present, the most effective available metric is employed to assess
 475 the impact of the proposed technique, with a primary focus on illustrating its effects in this particular
 476 context.

477

478

479 Societal impact Mirage Atom Diffusion advances de novo crystalline material generation by letting
 480 diffusion models vary atom counts during synthesis, boosting diversity and quality and accelerating
 481 discoveries for clean energy, electronics, and medicine. Its compatibility with multiple generative
 482 frameworks broadens access. But greater flexibility brings risks: potential design of hazardous
 483 or harmful materials without strong oversight, bias from narrow metrics like S.U.N., and high
 484 environmental and financial costs to validate many candidates. Growing model complexity also
 485 heightens concerns over interpretability, reproducibility, and equitable compute access.

483

484

485

486 REFERENCES
487

488 Mila AI4Science, Alex Hernandez-Garcia, Alexandre Duval, Alexandra Volokhova, Yoshua Bengio,
489 Divya Sharma, Pierre Luc Carrier, Yasmine Benabed, Michał Koziarski, and Victor Schmidt.
490 Crystal-gfn: sampling crystals with desirable properties and constraints, 2023. URL <https://arxiv.org/abs/2310.04925>.
491

492 Luis M Antunes, Keith T Butler, and Ricardo Grau-Crespo. Crystal structure generation with
493 autoregressive large language modeling. *Nature Communications*, 15(1):1–16, 2024.

494 Jacob Austin, Daniel D. Johnson, Jonathan Ho, Daniel Tarlow, and Rianne van den Berg. Structured
495 denoising diffusion models in discrete state-spaces, 2023. URL <https://arxiv.org/abs/2107.03006>.
496

497 Luis Barroso-Luque, Muhammed Shuaibi, Xiang Fu, Brandon M Wood, Misko Dzamba, Meng
498 Gao, Ammar Rizvi, C Lawrence Zitnick, and Zachary W Ulissi. Open materials 2024 (omat24)
499 inorganic materials dataset and models. *arXiv preprint arXiv:2410.12771*, 2024.
500

501 Zhendong Cao, Xiaoshan Luo, Jian Lv, and Lei Wang. Space group informed transformer for
502 crystalline materials generation. *arXiv preprint arXiv:2403.15734*, 2024.

503 Ricky TQ Chen and Yaron Lipman. Riemannian flow matching on general geometries. *arXiv preprint
504 arXiv:2302.03660*, 2023.
505

506 Kishalay Das, Subhjoyoti Khastagir, Pawan Goyal, Seung-Cheol Lee, Satadeep Bhattacharjee, and
507 Niloy Ganguly. Periodic materials generation using text-guided joint diffusion model, 2025. URL
508 <https://arxiv.org/abs/2503.00522>.

509 Valentin De Bortoli, Emile Mathieu, Michael Hutchinson, James Thornton, Yee Whye Teh, and
510 Arnaud Doucet. Riemannian score-based generative modeling. *arXiv preprint arXiv:2202.02763*,
511 2022.
512

513 Bowen Deng, Peichen Zhong, KyuJung Jun, Janosh Riebesell, Kevin Han, Christopher J Bartel, and
514 Gerbrand Ceder. Chgnet as a pretrained universal neural network potential for charge-informed
515 atomistic modelling. *Nature Machine Intelligence*, 5(9):1031–1041, 2023.

516 Bowen Deng, Yunyeong Choi, Peichen Zhong, Janosh Riebesell, Shashwat Anand, Zhuohan Li,
517 KyuJung Jun, Kristin A Persson, and Gerbrand Ceder. Systematic softening in universal machine
518 learning interatomic potentials. *npj Computational Materials*, 11(1):1–9, 2025.

519 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep
520 bidirectional transformers for language understanding, 2019. URL <https://arxiv.org/abs/1810.04805>.
521

522 Alex M. Ganose, Hrushikesh Sahasrabuddhe, Mark Asta, Kevin Beck, Tathagata Biswas, Alexander
523 Bonkowski, Joana Bustamante, Xin Chen, Yuan Chiang, Daryl Chrzan, Jacob Clary, Orion
524 Cohen, Christina Ertural, Max Gallant, Janine George, Sophie Gerits, Rhys Goodall, Rishabh
525 Guha, Geoffroy Hautier, Matthew Horton, Aaron Kaplan, Ryan Kingsbury, Matthew Kuner,
526 Bryant Li, Xavier Linn, Matthew McDermott, Rohith Srinivas Mohanakrishnan, Aakash Naik,
527 Jeffrey Neaton, Kristin Persson, Guido Petretto, Thomas Purcell, Francesco Ricci, Benjamin Rich,
528 Janosh Riebesell, Gian-Marco Rignanese, Andrew Rosen, Matthias Scheffler, Jonathan Schmidt,
529 Jimmy-Xuan Shen, Andrei Sobolev, Ravishankar Sundararaman, Cooper Tezak, Victor Trinquet,
530 Joel Varley, Derek Vigil-Fowler, Duo Wang, David Waroquiers, Mingjian Wen, Han Yang, Hui
531 Zheng, Jiongzh Zheng, Zhuoying Zhu, and Anubhav Jain. Atomate2: Modular workflows for
532 materials science. *ChemRxiv*, 2025. URL <https://chemrxiv.org/engage/chemrxiv/article-details/678e76a16dde43c9085c75e9>.
533

534 Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair,
535 Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information
536 processing systems*, 27, 2014.
537

538 Nate Gruver, Anuroop Sriram, Andrea Madotto, Andrew Gordon Wilson, C Lawrence Zitnick, and
539 Zachary Ulissi. Fine-tuned language models generate stable inorganic materials as text. *arXiv
preprint arXiv:2402.04379*, 2024.

540 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. [arXiv preprint](https://arxiv.org/abs/2006.11239)
 541 [arXiv:2006.11239](https://arxiv.org/abs/2006.11239), 2020.

542

543 Emiel Hoogeboom, Didrik Nielsen, Priyank Jaini, Patrick Forré, and Max Welling. Argmax flows
 544 and multinomial diffusion: Learning categorical distributions. [Advances in Neural Information](https://paperswithcode.com/paper/argmax-flows-and-multinomial-diffusion-learning-categorical-distributions)
 545 [Processing Systems](https://paperswithcode.com/paper/argmax-flows-and-multinomial-diffusion-learning-categorical-distributions), 34:12454–12465, 2021.

546 Emiel Hoogeboom, Victor Garcia Satorras, Clément Vignac, and Max Welling. Equivariant diffusion
 547 for molecule generation in 3d, 2022. URL <https://arxiv.org/abs/2203.17003>.

548

549 Chin-Wei Huang, Milad Aghajohari, Joey Bose, Prakash Panangaden, and Aaron C Courville.
 550 Riemannian diffusion models. [Advances in Neural Information Processing Systems](https://paperswithcode.com/paper/riemannian-diffusion-models), 35:2750–
 551 2761, 2022.

552 Anubhav Jain, Shyue Ping Ong, Geoffroy Hautier, Wei Chen, William Davidson Richards, Stephen
 553 Dacek, Shreyas Cholia, Dan Gunter, David Skinner, Gerbrand Ceder, and Kristin A. Persson.
 554 Commentary: The materials project: A materials genome approach to accelerating materials
 555 innovation. [APL Materials](https://doi.org/10.1063/1.4812323), 1(1):011002, 07 2013a. ISSN 2166-532X. doi: 10.1063/1.4812323.
 556 URL <https://doi.org/10.1063/1.4812323>.

557

558 Anubhav Jain, Shyue Ping Ong, Geoffroy Hautier, Wei Chen, William Davidson Richards, Stephen
 559 Dacek, Shreyas Cholia, Dan Gunter, David Skinner, Gerbrand Ceder, and Kristin a. Persson.
 560 The Materials Project: A materials genome approach to accelerating materials innovation. [APL](https://doi.org/10.1063/1.4812323)
 561 [Materials](https://doi.org/10.1063/1.4812323), 1(1):011002, 2013b. ISSN 2166532X. doi: 10.1063/1.4812323. URL [http://link.aip.org/link/AMPADS/v1/i1/p011002/s1&Agg=doi](https://aip.org/link/AMPADS/v1/i1/p011002/s1&Agg=doi).

562

563 Rui Jiao, Wenbing Huang, Peijia Lin, Jiaqi Han, Pin Chen, Yutong Lu, and Yang Liu. Crystal
 564 structure prediction by joint equivariant diffusion, 2024a. URL <https://arxiv.org/abs/2309.04475>.

565

566 Rui Jiao, Wenbing Huang, Yu Liu, Deli Zhao, and Yang Liu. Space group constrained crystal
 567 generation, 2024b. URL <https://arxiv.org/abs/2402.03992>.

568

569 Chaitanya K. Joshi, Xiang Fu, Yi-Lun Liao, Vahe Gharakhanyan, Benjamin Kurt Miller, Anuroop
 570 Sriram, and Zachary W. Ulissi. All-atom diffusion transformers: Unified generative modelling of
 571 molecules and materials, 2025. URL <https://arxiv.org/abs/2503.03965>.

572

573 J. Jumper, R. Evans, and A. Pritzel. Highly accurate protein structure prediction with alphafold, 2021.
 574 URL <https://doi.org/10.1038/s41586-021-03819-2>.

575

576 Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative
 577 adversarial networks. In [Proceedings of the IEEE/CVF conference on computer vision and pattern recognition](https://paperswithcode.com/paper/tero-karras-a-style-based-generator-architecture-for-generative-adversarial-networks), pp. 4401–4410, 2019.

578

579 Nikita Kazeev, Wei Nong, Ignat Romanov, Ruiming Zhu, Andrey Ustyuzhanin, Shuya Yamazaki,
 580 and Kedar Hippalgaonkar. Wyckoff transformer: Generation of symmetric crystals, 2025. URL
 581 <https://arxiv.org/abs/2503.02407>.

582

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

586

587 Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017. URL
 588 <https://arxiv.org/abs/1412.6980>.

589

590 Diederik P Kingma and Max Welling. Auto-encoding variational bayes. [arXiv preprint](https://arxiv.org/abs/1312.6114)
 591 [arXiv:1312.6114](https://arxiv.org/abs/1312.6114), 2013.

592

593 Leon Klein, Andreas Krämer, and Frank Noé. Equivariant flow matching, 2023. URL <https://arxiv.org/abs/2306.15030>.

594

595 Walter Kohn and Lu Jeu Sham. Self-consistent equations including exchange and correlation effects.
 596 [Physical review](https://doi.org/10.1103/PhysRev.140(4A):A1133), 140(4A):A1133, 1965.

594 G. Kresse and J. Furthmüller. Efficient iterative schemes for ab initio total-energy calculations using a
595 plane-wave basis set. *Phys. Rev. B*, 54:11169–11186, Oct 1996. doi: 10.1103/PhysRevB.54.11169.
596 URL <https://link.aps.org/doi/10.1103/PhysRevB.54.11169>.

597

598 Daniel Levy, Siba Smarak Panigrahi, Sékou-Oumar Kaba, Qiang Zhu, Kin Long Kelvin Lee, Mikhail
599 Galkin, Santiago Miret, and Siamak Ravankhahsh. Symmcmd: Symmetry-preserving crystal
600 generation with diffusion models, 2025. URL <https://arxiv.org/abs/2502.03638>.

601 Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching
602 for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.

603

604 Shuqi Lu, Huawei Lin, Lin Yao, Zhifeng Gao, Xiaohong Ji, Weinan E, Linfeng Zhang, and Guolin
605 Ke. Uni-3dar: Unified 3d generation and understanding via autoregression on compressed spatial
606 tokens, 2025. URL <https://arxiv.org/abs/2503.16278>.

607 Xiaoshan Luo, Zhenyu Wang, Qingchang Wang, Jian Lv, Lei Wang, Yanchao Wang, and Yanming
608 Ma. Crystalfow: A flow-based generative model for crystalline materials, 2025. URL <https://arxiv.org/abs/2412.11693>.

609

610 Benjamin Kurt Miller, Ricky T. Q. Chen, Anuroop Sriram, and Brandon M Wood. Flowmm:
611 Generating materials with riemannian flow matching, 2024. URL <https://arxiv.org/abs/2406.04713>.

612

613 Tri Minh Nguyen, Sherif Abdulkader Tawfik, Truyen Tran, Sunil Gupta, Santu Rana, and Svetha
614 Venkatesh. Efficient symmetry-aware materials generation via hierarchical generative flow net-
615 works, 2024. URL <https://arxiv.org/abs/2411.04323>.

616

617 Andrey Okhotin, Dmitry Molchanov, Vladimir Arkhipkin, Grigory Bartosh, Viktor Ohanesian, Aibek
618 Alanov, and Dmitry Vetrov. Star-shaped denoising diffusion probabilistic models, 2023. URL
619 <https://arxiv.org/abs/2302.05259>.

620

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

625

626 Janosh Riebesell, Rhys EA Goodall, Philipp Benner, Yuan Chiang, Bowen Deng, Alpha A Lee,
627 Anubhav Jain, and Kristin A Persson. Matbench Discovery—a framework to evaluate machine
628 learning crystal stability predictions. *arXiv preprint arXiv:2308.14920*, 2023.

629

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

632

633 Arne Schneuing, Ilia Igashov, Adrian W. Doppelstein, Thomas Castiglione, Michael M. Bronstein,
634 and Bruno Correia. Multi-domain distribution learning for de novo drug design. In *The Thirteenth
635 International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=g3VCIM94ke>.

636

637 Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised
638 learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*,
639 pp. 2256–2265. PMLR, 2015.

640

641 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben
642 Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint
643 arXiv:2011.13456*, 2020.

644

645 Anuroop Sriram, Benjamin Kurt Miller, Ricky T. Q. Chen, and Brandon M. Wood. Flowllm: Flow
646 matching for material generation with large language models as base distributions, 2024. URL
647 <https://arxiv.org/abs/2410.23405>.

648

649 Wenhao Sun, Stephen T Dacek, Shyue Ping Ong, Geoffroy Hautier, Anubhav Jain, William D
650 Richards, Anthony C Gamst, Kristin A Persson, and Gerbrand Ceder. The thermodynamic scale of
651 inorganic crystalline metastability. *Science advances*, 2(11):e1600225, 2016.

648 Kasper Tolborg, Johan Klarbring, Alex M Ganose, and Aron Walsh. Free energy predictions for
649 crystal stability and synthesisability. *Digital Discovery*, 1(5):586–595, 2022.
650

651 Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée
652 Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand
653 Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation language
654 models, 2023. URL <https://arxiv.org/abs/2302.13971>.

655 Arash Vahdat, Karsten Kreis, and Jan Kautz. Score-based generative modeling in latent space.
656 *Advances in Neural Information Processing Systems*, 34, 2021.
657

658 Hanlin Wu, Yuxuan Song, Jingjing Gong, Ziyao Cao, Yawen Ouyang, Jianbing Zhang, Hao Zhou,
659 Wei-Ying Ma, and Jingjing Liu. A periodic bayesian flow for material generation, 2025. URL
660 <https://arxiv.org/abs/2502.02016>.

661 Zhisheng Xiao, Karsten Kreis, and Arash Vahdat. Tackling the generative learning trilemma with
662 denoising diffusion gans, 2022. URL <https://arxiv.org/abs/2112.07804>.
663

664 Tian Xie, Xiang Fu, Octavian-Eugen Ganea, Regina Barzilay, and Tommi Jaakkola. Crystal diffusion
665 variational autoencoder for periodic material generation, 2022. URL <https://arxiv.org/abs/2110.06197>.
666

667 Sherry Yang, KwangHwan Cho, Amil Merchant, Pieter Abbeel, Dale Schuurmans, Igor Mordatch,
668 and Ekin Dogus Cubuk. Scalable diffusion for materials generation, 2024. URL <https://arxiv.org/abs/2311.09235>.
669

670 Claudio Zeni, Robert Pinsler, Daniel Zügner, Andrew Fowler, Matthew Horton, Xiang Fu, Sasha
671 Shysheya, Jonathan Crabbé, Lixin Sun, Jake Smith, Bichlien Nguyen, Hannes Schulz, Sarah
672 Lewis, Chin-Wei Huang, Ziheng Lu, Yichi Zhou, Han Yang, Hongxia Hao, Jielan Li, Ryota
673 Tomioka, and Tian Xie. MatterGen: a generative model for inorganic materials design, 2024. URL
674 <https://arxiv.org/abs/2312.03687>.
675

676

677

678

679

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702 **A CRYSTAL INVARIANCES**
703704 **A.1 SYMMETRIES OF CRYSTAL STRUCTURE DISTRIBUTION**
705706 The space of 3D crystal structures is governed by symmetries that impose stringent constraints
707 and play a crucial role in directing the generation process from the distribution $p(\mathcal{M})$. The model,
708 proposed in Jiao et al. (2024a) and described in Section 2.3, focuses on the following symmetries:
709710 **Permutation invariance:** \forall permutation $P \in S_N \implies p(L, FP, AP) = p(L, F, A)$
711712 **$O(3)$ invariance:** \forall orthogonal transformation $Q \in \mathbb{R}^{3 \times 3} \implies p(QL, F, A) = p(L, F, A)$
713714 **Periodic translation invariance:** \forall translation $\tau \in \mathbb{R}^{1 \times 3} \implies p(L, w(F + \mathbf{1}\tau), A) = p(L, F, A)$
715716 Where $w(F) = F - \lfloor F \rfloor \in [0, 1)^{N_{\text{atoms}} \times 3}$ returns the fractional part of each element in F and
717 $\mathbf{1} \in \mathbb{R}^{N_{\text{atoms}} \times 1}$ is a unit vector.
718719 Informally, each symmetry signifies that performing a particular transformation on a crystal does not
720 alter its likelihood under the distribution $p(\mathcal{M})$. These properties provide a strong inductive bias and
721 can be methodically integrated into the diffusion model by design.
722723 **A.2 INVARIANT DIFFUSION MODELS**
724725 A diffusion model $p_\theta(x_0)$ is said to be invariant with respect to the symmetry group \mathbb{G} if, for any
726 transformation $g \in \mathbb{G}$, it holds that $p_\theta(g \cdot x_0) = p_\theta(x_0)$. Hoogeboom et al. (2022) proposed that
727 invariance can be encapsulated in the diffusion model if we define:
728729 **Invariant prior distribution:** $\forall g \in \mathbb{G} \implies p(x_T) = p(g \cdot x_T)$ (3)
730731 **Equivariant transition kernels:** $\forall g \in \mathbb{G} \implies p_\theta(g \cdot x_{t-1} | g \cdot x_t) = p_\theta(x_{t-1} | x_t)$ (4)
732733 $q(g \cdot x_t | g \cdot x_{t-1}) = q(x_t | x_{t-1})$ (5)734 If conditions 3, 4 hold, then the diffusion model $p_\theta(x_0)$ is invariant with respect to the symmetry
735 group \mathbb{G} :
736

737
$$p_\theta(g \cdot x_0) = \int p(g \cdot x_T) \prod_{t=1}^T p_\theta(g \cdot x_{t-1} | g \cdot x_t) dx_{1:T} = \int p(x_T) \prod_{t=1}^T p_\theta(x_{t-1} | x_t) dx_{1:T} = p_\theta(x_0),$$

738

739 whereas if conditions 4, 5 hold, then the conditional backward process is equivariant:
740

741
$$\begin{aligned} q(g \cdot x_{t-1} | g \cdot x_t, g \cdot x_0) &= \frac{q(g \cdot x_t | g \cdot x_{t-1}) q(g \cdot x_{t-1} | g \cdot x_0)}{q(g \cdot x_t | g \cdot x_0)} = \\ &= \frac{q(x_t | x_{t-1}) q(x_{t-1} | x_0)}{q(x_t | x_0)} = q(x_{t-1} | x_t, x_0), \end{aligned}$$

742

743 and the training objective is invariant:
744

745
$$\begin{aligned} \mathcal{L}_g &= \sum_{t=2}^T \gamma_t \mathbb{E}_{g \cdot x_0 \sim q(g \cdot x_0), g \cdot x_t \sim q(g \cdot x_t | g \cdot x_0)} D_{\text{KL}} [q(g \cdot x_{t-1} | g \cdot x_t, g \cdot x_0) \parallel p_\theta(g \cdot x_{t-1} | g \cdot x_t)] = \\ &= \sum_{t=2}^T \gamma_t \mathbb{E}_{x_0 \sim q(x_0), x_t \sim q(x_t | x_0)} D_{\text{KL}} [q(x_{t-1} | x_t, x_0) \parallel p_\theta(x_{t-1} | x_t)] = \mathcal{L} \end{aligned}$$

746

747 The same remains true for score-matching objective if in addition to the condition 5 the following
748 conditions are also satisfied: 6 – equivariance of the score-estimator $s_\theta(x_t, t)$, 7 – distributivity of the
749 transformations g from the group \mathbb{G} , 8 – invariance of l_2 -norm:
750

751
$$s_\theta(g \cdot x_t, t) = g \cdot s_\theta(x_t, t) \tag{6}$$

752

753
$$g \cdot x + g \cdot y = g \cdot (x + y) \tag{7}$$

754

755
$$\|x\|_2^2 = \|g \cdot x\|_2^2 \tag{8}$$

756 Then the score-matching objective is invariant with respect to the symmetry group \mathbb{G} :

$$\begin{aligned}
 \mathcal{L}_g &= \mathbb{E}_{g \cdot x_0 \sim q(g \cdot x_0), t \sim \mathcal{U}(1, T), g \cdot x_t \sim q(g \cdot x_t | g \cdot x_0)} \|\nabla_{g \cdot x_t} \log p(g \cdot x_t | g \cdot x_0) - s_\theta(g \cdot x_t, t)\|_2^2 = \\
 &= \mathbb{E}_{x_0 \sim q(x_0), t \sim \mathcal{U}(1, T), x_t \sim q(x_t | x_0)} \|g \cdot \nabla_{x_t} \log p(x_t | x_0) - g \cdot s_\theta(x_t, t)\|_2^2 = \\
 &= \mathbb{E}_{x_0 \sim q(x_0), t \sim \mathcal{U}(1, T), x_t \sim q(x_t | x_0)} \|g \cdot (\nabla_{x_t} \log p(x_t | x_0) - s_\theta(x_t, t))\|_2^2 = \\
 &= \mathbb{E}_{x_0 \sim q(x_0), t \sim \mathcal{U}(1, T), x_t \sim q(x_t | x_0)} \|\nabla_{x_t} \log p(x_t | x_0) - s_\theta(x_t, t)\|_2^2 = \mathcal{L}
 \end{aligned}$$

763 The definition of the diffusion model in terms of the score function allows for the formulation of
764 various backward transition kernels, denoted as K_θ :

$$x_{t-1} = K_\theta(x_t)$$

766 These kernels incorporate the score estimator s_θ in a specific manner and can be either deterministic,
767 or stochastic. Equivariance with respect to the symmetry group \mathbb{G} , in the context of transition kernels
768 K , implies the following condition in the deterministic case:

$$\forall g \in \mathbb{G}, \quad K(g \cdot x) = g \cdot K(x) \quad (9)$$

771 In the stochastic case, this equality is sufficient. However, the necessary and sufficient condition
772 requires equality solely in terms of probability density functions.

773 A.2.1 INVARIANCES

775 Jiao et al. (2024a) proposed that symmetric properties outlined in Section A.1 can be encapsulated in
776 the joint diffusion model described in Section 2.3 via specific neural network parameterization.

778 **Permutation invariance** affects the components F and A . The prior distributions for these
779 components are defined element-wise, and the corresponding transition kernels remain equivariant
780 as long as the neural network is equivariant. It is achieved using a graph neural network (GNN)
781 architecture that alternates between message-passing and atom-wise processing layers. The message-
782 passing layer operates on atom pairs selected based on rules that are independent of their position
783 in the sequence, i.e., distance cutoff between atoms. As a result, reordering of the input sequence
784 of atoms results in the same reordering of the output sequence of predictions, ensuring permutation
785 equivariance.

786 **$O(3)$ invariance** affects only the L component of the crystal because the F component defines
787 coordinates within the unit cell and does not contain any information about the orientation of the unit
788 cell in space. The prior distribution in DDPM is already invariant. In order to define an equivariant
789 transition kernel we need to parametrize its mean in the following way:

$$\mu_\theta(\mathcal{M}_t, t) = \mu_\theta(L_t, F_t, A_t, t) = L_t \text{NN}_\theta(L_t^T L_t, F_t, A_t, t),$$

790 where $\text{NN}_\theta : \mathcal{M}_t, t \rightarrow \mathbb{R}^{3 \times 3}$ – neural network with a linear layer on the top. Then, the following
791 transition kernel in DDPM:

$$p_\theta(L_{t-1} | \mathcal{M}_t) = p_\theta(L_{t-1} | L_t, F_t, A_t) = \mathcal{N}(L_{t-1} | \mu_\theta(L_t, F_t, A_t, t), \sigma_t I)$$

792 is $O(3)$ equivariant:

$$\begin{aligned}
 p_\theta(QL_{t-1} | QL_t, F_t, A_t) &= \mathcal{N}(QL_{t-1} | \mu_\theta(QL_t, F_t, A_t, t), \sigma_t I) = \\
 &= \mathcal{N}(QL_{t-1} | QL_t \text{NN}_\theta(L_t^T Q^T Q L_t, F_t, A_t, t), \sigma_t I) = \\
 &= \mathcal{N}(QL_{t-1} | QL_t \text{NN}_\theta(L_t^T L_t, F_t, A_t, t), \sigma_t I) = \\
 &= \mathcal{N}(QL_{t-1} | Q\mu_\theta(L_t, F_t, A_t, t), \sigma_t I) = \\
 &= \mathcal{N}(L_{t-1} | \mu_\theta(L_t, F_t, A_t, t), \sigma_t I) = \\
 &= p_\theta(L_{t-1} | L_t, F_t, A_t)
 \end{aligned}$$

804 **Periodic translation invariance** affects only the F component. The prior distribution in Wrapped
805 Normal diffusion is already invariant. To ensure that the transition kernel is equivariant, we parame-
806 terize the score estimator s_θ in an invariant form:

$$\begin{aligned}
 s_\theta(\mathcal{M}_t, t) &= s_\theta(L_t, F_t, A_t, t) = s_\theta\left(L_t, \text{PairwiseDist}(F_t), A_t, t\right), \\
 \text{PairwiseDist}(F_t) &= \left\{ \psi_{FT}(f_j - f_i) \mid i, j = 1, \overline{N_{\text{atoms}}}, i \neq j \right\},
 \end{aligned}$$

810 where f_i – fractional coordinates of the atom i from the F_t , $\psi_{FT} : (-1, 1)^3 \rightarrow [-1, 1]^{3 \times K}$
 811 – the Fourier Transformation of the relative fractional coordinate $f_i - f_j$, and $\text{PairwiseDist} : [0, 1]^{N_{\text{atoms}} \times 3} \rightarrow [-1, 1]^{N_{\text{atoms}} \times N_{\text{atoms}} \times K}$ – yields coordinates representation which is invariant to
 812 periodic translations:
 813

$$\text{PairwiseDist}(w(F_t + \mathbf{1}\tau)) = \text{PairwiseDist}(F_t)$$

816 This follows from the fact that pairwise atomic distances remain unchanged when the entire system
 817 of atoms is translated. Further, if we use the following form of the stochastic backward transition
 818 kernel K_θ :

$$F_{t-1} = K_\theta(\mathcal{M}_t) = K_\theta(L_t, F_t, A_t) = w(a_t F_t + b_t s_\theta(L_t, F_t, A_t, t) + c_t \epsilon), \quad \epsilon \sim \mathcal{N}(0, 1),$$

820 where a_t, b_t, c_t are scalar coefficients, then K_θ is periodic translation equivariant (sufficient condition
 821 9):
 822

$$\begin{aligned} K_\theta(L_t, w(F_t + \mathbf{1}\tau), A_t) &= w(a_t w(F_t + \mathbf{1}\tau) + b_t s_\theta(L_t, w(F_t + \mathbf{1}\tau), A_t, t) + c_t \epsilon) = \\ &= w(a_t w(F_t + \mathbf{1}\tau) + b_t s_\theta(L_t, F_t, A_t, t) + c_t \epsilon) = \\ &= w(a_t F_t + b_t s_\theta(L_t, F_t, A_t, t) + c_t \epsilon + \mathbf{1}\tau) = \\ &= w(w(a_t F_t + b_t s_\theta(L_t, F_t, A_t, t) + c_t \epsilon) + \mathbf{1}\tau) = \\ &= w(K_\theta(L_t, F_t, A_t) + \mathbf{1}\tau) \end{aligned}$$

B ABLATION

832 As discussed in Section 6, several critical steps were undertaken in developing the final version of the
 833 MiAD.

B.1 NUMBER OF MIRAGE ATOMS

836 The definition of mirage infusion necessitates setting the hyperparameter N_m , which specifies that
 837 crystals are supplemented with mirage atoms until the total number of atoms reaches N_m . Increasing
 838 this hyperparameter results in a greater number of possible atom variants from which the model can
 839 select during crystal construction. Simultaneously, it increases the number of atoms the model must
 840 eliminate during the generation process. The neural network architecture proposed by Jiao et al.
 841 (2024a) has a limit on the number of atoms that can interact efficiently. Beyond a certain number, the
 842 neural network loses the capacity to manage them effectively. We evaluate several variants of N_m
 843 in Table 3 using S.U.N. computed via MLIPs (see Section 5), and select the optimal variant for all
 844 subsequent experiments.

B.2 LOSS COMPONENT PRIORITIZATION

847 The application of mirage infusion alters the loss components associated with diffusion in fractional
 848 coordinates and atom types. This, in turn, affects the scale of the gradients and the prioritization
 849 of tasks the neural network must solve concurrently: 1) lattice prediction, 2) fractional coordinates
 850 prediction, and 3) atom types prediction. We discovered that the balance among these tasks, espe-
 851 cially the influence of the loss component associated with atom types \mathcal{L}_A (see Section 2.3.1), can
 852

854 **Table 3: Ablation study on the number of mirage atoms in MiAD** We perform a comparison
 855 of various values for N_m within MiAD: 20, 25, 30, 35. This hyperparameter specifies that crystals
 856 are supplemented with mirage atoms until the total number of atoms reaches N_m . The models are
 857 compared using S.U.N. for 10 000 sampled crystals, where stability is estimated via (left) eq-V2 and
 858 (right) CHGNet.

Model	eq-V2 ($E^{\text{hull}} < 0.0$)				CHGNet ($E^{\text{hull}} < 0.0$)			
	Stable (%) ↑	Unique (%) ↑	Novel (%) ↑	S.U.N. (%) ↑	Stable (%) ↑	Unique (%) ↑	Novel (%) ↑	S.U.N. (%) ↑
MiAD (20)	8.7	92.0	73.0	5.3	17.4	91.9	72.9	12.0
MiAD (25)	9.7	92.2	71.1	5.5	19.8	92.2	71.3	12.9
MiAD (30)	8.3	93.1	74.1	4.7	16.6	93.2	73.9	11.3
MiAD (35)	8.2	92.3	70.9	4.6	17.9	92.4	70.9	11.8

864
865 **Table 4: Ablation study of loss scaling for atom types in MiAD** We perform a comparison of
866 the coefficients 0.5, 1.0, 2.0 applied to the loss function for atom types in MiAD, while maintaining
867 constant scales for the losses associated with lattice and fractional coordinates. We quantify the
868 prioritization of the loss components corresponding to $(L - F - A)$ at the end of the training procedure
869 as a percentage of the total loss. The prioritization for these coefficients are as follows: $\mathcal{L}_A \times 0.5$:
870 $(40 - 50 - 10)$, $\mathcal{L}_A \times 1.0$: $(36 - 46 - 18)$, $\mathcal{L}_A \times 2.0$: $(31 - 39 - 30)$. The models are compared
871 using S.U.N. for 10 000 sampled crystals, where stability is estimated via (left) eq-V2 and (right)
872 CHGNet.
873

Model	eq-V2 ($E^{\text{hull}} < 0.0$)				CHGNet ($E^{\text{hull}} < 0.0$)			
	Stable (%) ↑	Unique (%) ↑	Novel (%) ↑	S.U.N. (%) ↑	Stable (%) ↑	Unique (%) ↑	Novel (%) ↑	S.U.N. (%) ↑
MiAD ($\mathcal{L}_A \times 0.5$)	8.2	93.8	77.6	4.7	17.9	93.8	77.6	12.3
MiAD ($\mathcal{L}_A \times 1.0$)	9.7	92.2	71.1	5.5	19.8	92.2	71.3	12.9
MiAD ($\mathcal{L}_A \times 2.0$)	8.9	91.5	67.4	5.0	19.0	91.5	67.4	11.9

874
875 **Table 5: Ablation study of possible definitions of the mirage infusion** We perform a comparison
876 of the definitions of MiAD in terms of (1) the initialization of fractional coordinates of mirage atoms
877 in crystals: sampling from the uniform distribution or positioning at the geometric center of mass of
878 the real atoms, (2) the masking of mirage atoms in the loss for fractional coordinates. The models are
879 compared using S.U.N. for 10 000 sampled crystals, where stability is estimated via (left) eq-V2 and
880 (right) CHGNet.
881

Model	eq-V2 ($E^{\text{hull}} < 0.0$)				CHGNet ($E^{\text{hull}} < 0.0$)			
	Stable (%) ↑	Unique (%) ↑	Novel (%) ↑	S.U.N. (%) ↑	Stable (%) ↑	Unique (%) ↑	Novel (%) ↑	S.U.N. (%) ↑
MiAD (Uniform + Masked)	9.7	92.2	71.1	5.5	19.8	92.2	71.3	12.9
MiAD (Uniform + NonMasked)	7.7	91.7	65.1	3.8	16.6	91.6	65.2	10.3
MiAD (Center + Masked)	0.5	93.4	93.1	0.4	5.8	92.9	93.0	5.2
MiAD (Center + NonMasked)	2.6	96.3	87.9	1.3	9.9	96.3	87.9	8.1

882 substantially affect the model’s quality. During our experiments, we did not modify the scales of
883 the loss components related to lattice and fractional coordinates in order to clearly demonstrate the
884 impact of the proposed technique. However, the loss associated with atom types is significantly
885 influenced by the hyperparameter N_m in mirage infusion, which raises the question of whether
886 additional corrections to this component are necessary. We quantify the prioritization of the loss
887 components corresponding to $(L - F - A)$ at the end of the training procedure as a percentage of the
888 total loss. The application of mirage infusion with $N_m = 25$ reduces by half the loss for atom types,
889 denoted as \mathcal{L}_A , and leads to a balance of $(36 - 46 - 18)$ among MiAD components, whereas the
890 balance among components in the original DiffCSP is $(31 - 39 - 30)$. Then, in Table 4, we compare
891 MiAD models, where \mathcal{L}_A is further increased or decreased by a factor of two, however, it only
892 diminishes the quality. We considered it essential to illustrate the effects of these adjustments, due to
893 the scale of an impact on the model’s quality. In Table 4, the impact of \mathcal{L}_A as a percentage of the total
894 loss is not precisely doubled or halved because alterations to this component also influence the total
895 loss. Given the significance of this loss prioritization, we conducted experiments, the results of which
896 are presented in Table 3, using coefficients for \mathcal{L}_A that maintain the same balance of $(36 - 46 - 18)$
897 among the loss components by the end of the training procedure.
898

900 B.3 DEFINITIONS OF MIRAGE INFUSION

901 As discussed in Section 3, Schneuing et al. (2025) introduced a related concept involving the
902 augmentation of original molecular structures with hypothetical (non-existent) components for
903 structure-based drug design. In this approach, the fractional coordinates of mirage atoms are
904 initialized at the geometric center of mass of the real atoms, and no masking is applied to the mirage
905 atoms. Table 5 compares this formulation with the definition of mirage infusion proposed in the
906 present work, showing that the latter achieves substantially improved performance.
907

914 B.4 FINAL IMPACT OF MIRAGE INFUSION

915 Figure 3 illustrates the comparison between the final version of MiAD and the original DiffCSP. In
916 this analysis, we employed the optimal variant of $N_m = 25$ as identified in Table 3 and adhered to the
917

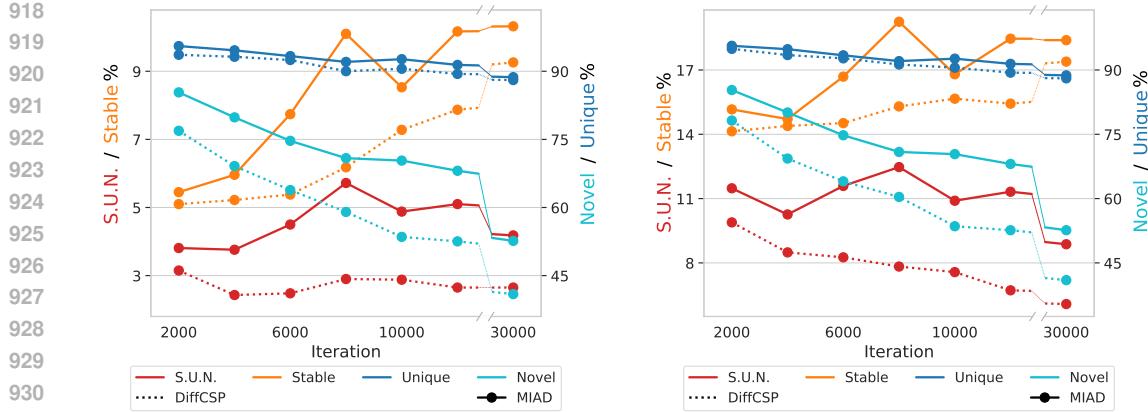


Figure 3: Comparison of MiAD (DiffCSP with mirage infusion) in its final version and DiffCSP in terms of stability, uniqueness, novelty, and S.U.N. Stability is estimated via (left) eq-V2 and (right) CHGNet. MiAD outperforms DiffCSP across all metrics, especially in terms of stability rate and S.U.N., achieving the highest quality after 8000 epochs.

default, yet optimal, prioritization of the loss components as specified in Table 4. The incorporation of mirage infusion notably enhances both stability and novelty; however, it also extends the time necessary to reach peak performance.

C EXPERIMENTAL DETAILS

Code A zip archive containing the raw code for all the experiments is accessible for download through this link.

Dataset All experiments were conducted using the MP-20 dataset (Jain et al., 2013a), which comprises 45,231 stable inorganic materials selected from Material Projects (Jain et al., 2013a). We employed the same train-validation-test split of 60-20-20 as used in the study by Xie et al. (2022).

Neural network architecture The models utilized the CSPNet neural network architecture proposed by Jiao et al. (2024a), with the following hyperparameters: "hidden_dim" : 512, "num_layers" : 6, "num_freqs" : 128, "latent_dim" : 256, "max_atoms" : 100, "act_fn" : "silu", "dis_emb" : "sin", "edge_style" : "fc", "max_neighbors" : 20, "cutoff" : 7.0, "ln" : True, "ip" : True. The application of mirage infusion, which necessitates the use of an additional atom type 0 in the D3PM diffusion module corresponding to atom types (refer to Sections 2.3.1 and 3), does not increase the number of neural network parameters. This is because the configuration proposed by Jiao et al. (2024a) already includes several free atom types that can be utilized in the proposed model.

Optimization Jiao et al. (2024a) proposed utilizing a batch size of 256 in conjunction with the Adam optimizer (Kingma & Ba, 2017) with an initial learning rate of 10^{-3} and a scheduler that reduces the learning rate to 10^{-4} when the validation loss ceases to decrease. We use the same batch size, but in contrast, we employed the Adam optimizer solely with a learning rate of 10^{-3} , omitting any schedulers. Our experiments indicated that the validation loss is not strongly correlated with the quality of the diffusion model once the model performance is sufficiently high. Additional experiments confirmed that this modification in the optimization procedure does not influence the quality of the default DiffCSP. We also found that 1000 epochs were adequate for training the original DiffCSP model, but insufficient for a model after the application of mirage infusion. For the best-proposed version of mirage infusion, maximum performance (as measured by S.U.N. computed via eq-V2) was achieved after 8000 epochs (see Figure 3). Due to this increased requirement for the number of epochs, we removed the scheduler to allow further improvements in the models after the 1000 epochs of training.

Computational costs The proposed mirage infusion technique (see Section 3) substantially increases the average number of atoms with which the neural network interacts during the training and sampling procedures, leading to increased time and memory costs. After applying the mirage infusion technique in its optimal configuration ($N_m = 25$, which corresponds to approximately a $\times 2.7$ increase in the average number of atoms in the crystals), we observe the following approximate increases in execution time: training step $\times 4.3$, sampling step $\times 3.8$; and the following approximate increases in memory consumption: training step $\times 4.4$, sampling step $\times 4.6$. These estimates were obtained on a single GPU NVIDIA Tesla V100 32 GB with 4 CPU cores Intel Xeon Gold 6152 (2.1–3.7 GHz). In our experiments, MiAD ($N_m = 25$) required 4 days for the training procedure (8000 epochs) and 2 hours for the sampling procedure (10 000 crystals) on 2 GPUs NVIDIA Tesla A100 80 GB with 8 CPU cores AMD EPYC 7702 2.3.35 HHZ. The training and sampling procedures can also be conducted on a single GPU, though this will incur increased time costs.

D DFT COMPUTATION DETAILS

We use DFT settings from Materials Project <https://docs.materialsproject.org/methodology/materials-methodology/calculation-details/gga+u-calculations/parameters-and-convergence> for structure relaxation and energy computation. In particular, we do GGA and GGA+U calculations with `atomate2.vasp.flows.mp`. `MPGGADoubleRelaxStaticMaker` (Ganose et al., 2025), which in turn relies on `pymatgen.io.vasp.sets.MPRelaxSet` and `pymatgen.io.vasp.sets.MPStaticSet` (Ong et al., 2013). Computations themselves were done with VASP (Kresse & Furthmüller, 1996) version 5.4.4. The raw total energies computed by DFT were corrected with MaterialsProject2020Compatibility before putting into the `PhaseDiagram` to obtain the DFT E^{hull} . We used the MP convex hull 2023-02-07-ppd-mp.pkl.gz distributed by `matbench-discovery` (Riebesell et al., 2023) as the reference hull.

E ADDITIONAL METRICS

While S.U.N. remains the principal criterion for evaluating modern generative models in de novo crystal generation Miller et al. (2024); Sriram et al. (2024); Joshi et al. (2025), it is instructive to consider additional metrics that are being used in the literature. These include Structure Validity, Compositional Validity, Coverage (COV-R, COV-P), and distributional distances such as the Wasserstein distance of scalar material properties (e.g., density, number of elements), proposed in Xie et al. (2022). Such metrics provide complementary perspectives on model performance, quantifying local consistency of atomic structures or alignment with empirical property distributions.

However, these measures exhibit important limitations. Particularly, Validity and Coverage are nearly saturated for modern models, as shown in Table 6, with reported differences between approaches often below 1%. More critically, they fail to penalize overfitting. For example, if the MP-20 training set is itself evaluated as a “generative model”, it achieves near-optimal values across Validity and Coverage, underscoring that such metrics can be artificially inflated just by replication of training data. Consequently, a model that prioritizes memorization over discovery may appear competitive according to these statistics, despite offering little value for materials discovery.

An illustrative case can be drawn by comparing MiAD with FlowLLM. FlowLLM reports a higher Compositional Validity (89.05% vs. 84.21%), yet its Uniqueness & Novelty (U&N) rate shown in Table 1 is markedly lower (33.8% vs. 65.2%). The discrepancy arises because FlowLLM predominantly reproduces training set compositions, which maximizes apparent validity while suppressing novelty. In contrast, MiAD sacrifices a small degree of compositional accuracy but generates a significantly larger proportion of genuinely new materials, which is the central objective of de novo generation.

To further validate the proposed MiAD approach, we conducted the experiments with small-scale datasets, Perov-5 and Carbon-24. For both of them, we employed the same mirage infusion configuration as for MP-20, specifically $N_m = N + 5$, where N is the maximum number of atoms in crystals in each respective dataset (i.e., $N_m = 10$ for Perov-5 and $N_m = 29$ for Carbon-24). All

1026
 1027 Table 6: Evaluation of MiAD on Perov-5, Carbon-24, and MP-20 datasets compared with baseline
 1028 models. All results follow the same mirage infusion configuration as for MP-20.

Data	Method	Validity (%)↑		Coverage (%)↑		Property (%)↓	
		Struc.	Comp.	COV-R	COV-P	d_ρ	d_{elem}
Perov-5	DiffCSP	100.00	98.85	99.74	98.27	0.111	0.013
	CrysBFN	100.00	98.86	99.52	98.63	0.073	0.010
	TGDMat	100.00	98.63	99.83	99.52	0.050	0.009
	MiAD	94.82	97.91	98.07	92.82	0.089	0.075
Carbon-24	DiffCSP	100.00	NA	99.90	97.27	0.081	NA
	CrysBFN	100.00	NA	99.90	99.12	0.061	NA
	TGDMat	100.00	NA	99.99	92.43	0.043	NA
	Uni-3DAR	99.99	NA	100.00	98.16	0.066	NA
	MiAD	99.85	NA	99.51	99.46	0.061	NA
MP-20	DiffCSP	100.00	83.25	99.71	99.76	0.350	0.340
	UniMat	97.20	89.40	99.80	99.70	0.088	0.056
	FlowMM	96.85	83.19	99.49	99.58	0.239	0.083
	FlowLLM	99.81	89.05	99.06	99.68	0.660	0.090
	SymmCD	90.34	85.81	99.58	97.76	0.230	0.400
	WyFormer+DiffCSP++	99.80	81.40	99.51	95.81	0.360	0.079
	SymmBFN	94.27	83.93	99.73	99.00	0.083	0.095
	CrysBFN	100.00	87.51	99.09	99.79	0.207	0.163
	TGDMat	100.00	92.97	99.89	99.95	0.338	0.289
	Uni-3DAR	99.89	90.31	99.62	99.83	0.477	0.069
	MiAD	99.25	84.21	99.35	99.80	0.233	0.027
	MP-20 Train	100.00	90.65	99.81	99.79	0.133	0.025

1053
 1054 other hyperparameters (batch size, number of epochs, and neural network architecture) were adopted
 1055 directly from DiffCSP.

1056
 1057 A comparison of models trained on the Perov-5 and Carbon-24 datasets using the S.U.N. metric is
 1058 not meaningful because both datasets contain materials that are thermodynamically unstable under
 1059 standard conditions Xie et al. (2022). Consequently, recent studies Zeni et al. (2024); Miller et al.
 1060 (2024); Sriram et al. (2024); Joshi et al. (2025); Kazeev et al. (2025) have discontinued the use of
 1061 these two datasets for benchmarking de novo generation. Currently, there are no well-established
 1062 methods for fair comparison on these datasets that penalize overfitting. Therefore, Validity, Coverage,
 1063 and Property could be used only to validate the model’s ability to generate coherent structures, while
 1064 ignoring overfitting. In this light, Table 6 shows that MiAD is competitive with the prior baselines.

1065
 1066 This analysis highlights why S.U.N. remains the most informative single metric. Unlike isolated
 1067 measures, S.U.N. directly evaluates the trade-off between replicating known structures and discovering
 1068 stable, previously unobserved crystals, thereby capturing the utility of generative models in the
 1069 materials discovery pipeline. While auxiliary metrics may still be valuable for diagnostic purposes,
 1070 they should be interpreted with caution, as they may obscure or even contradict the overarching goal
 1071 of novelty-driven generation.

1072 F DISTRIBUTION OF NUMBER OF ATOMS

1073
 1074 Mirage infusion incorporates the ability to insert and remove atoms during the generation process.
 1075 However, we must verify that the model actually learns to use this capability. To this end, we report
 1076 statistics on the number of atoms in S.U.N. crystals generated by MiAD (see Figure 4).

1077
 1078 $p(A_{T,i}) = \text{Cat}(A_{T,i} \mid \mathbf{1}/(N_{\text{types}} + 1))$ — atom types (including the mirage type) have a uniform prior
 1079 (see Section 2.3.1)), where $N_{\text{types}} = 100$ in practice. Thus, if we use mirage infusion with $N_m = 25$,
 then at the start of generation ($t = T$) all atoms in a crystal are real with probability ≈ 0.77 , exactly

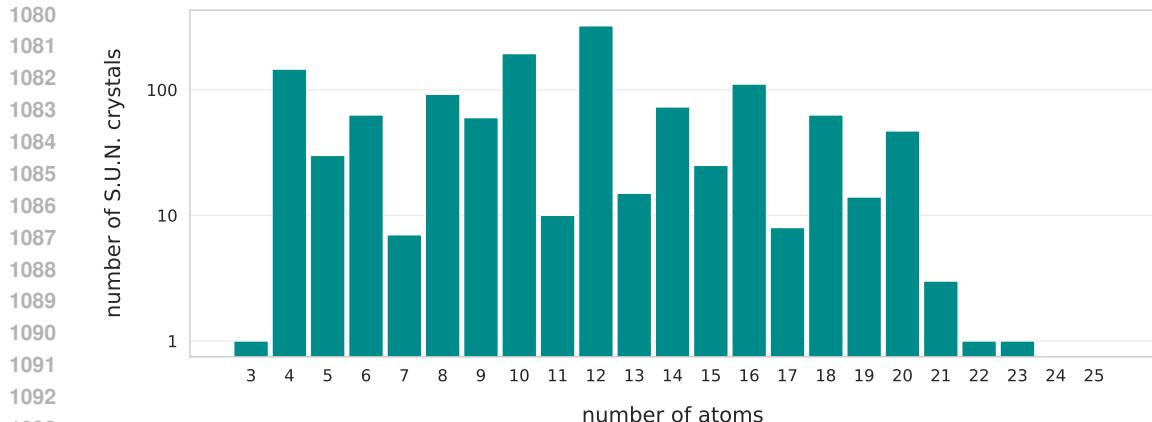


Figure 4: Number of atoms in S.U.N. crystals generated by MiAD. We consider only S.U.N. crystals among the 10 000 crystals generated by MiAD. After generation, crystals are prerelaxed via CHGNet. Stability is also estimated via CHGNet.

one atom is mirage with probability ≈ 0.19 , and so on. This implies that, at the start of generation, almost all crystals contain more than 23 real atoms. This fact, together with the statistics in Figure 4, demonstrates that MiAD

- changes the number of atoms in a crystal during generation;
- generates S.U.N. crystals with different numbers of atoms.

Crystals with 3, 7, 11, 13, 17 atoms appear more rarely than others, because crystals with these atom counts are underrepresented in the training data.

G SPACEGROUP DISTRIBUTION

Diversity is a key aspect of the quality of a generative model. S.U.N. addresses this aspect via U – uniqueness. However, this is not the only relevant evaluation, because crystals can differ from one another in various ways. At the same time, many existing diffusion-like models (as well as MiAD) for crystal generation cannot guarantee that the generated crystals will span different spacegroups (particularly the more complex ones, such as cubic or hexagonal). The common assumption is that a diffusion model will learn this from the training data.

To address these concerns, we compare the number of S.U.N. crystals belonging to each spacegroup among 10 000 crystals generated by different models and prerelaxed using CHGNet (see Figure 5). Based on this experiment, we can derive the following claims:

- Existing generative models for crystals do not suffer from severe mode collapse;
- MiAD, on average, generates more S.U.N. crystals than other models without exhibiting mode collapse.

From the application perspective, we argue that the particular shape of the spacegroup distribution is not one of the primary aspects of generative model quality in materials discovery. A spacegroup distribution that is closest to the train or test distribution, or that is the most uniform, does not by itself indicate clear benefits of a particular generative model for the task of inventing novel materials. At the same time, the shape of such distributions can indicate drawbacks if the generative model completely ignores crystals from a particular spacegroup. Our main point is that we should not conflate distribution-based quality criteria with S.U.N., but rather use

- S.U.N. as an averaged quality measure;
- distribution-based metrics as diagnostics to check for the absence of serious diversity issues.

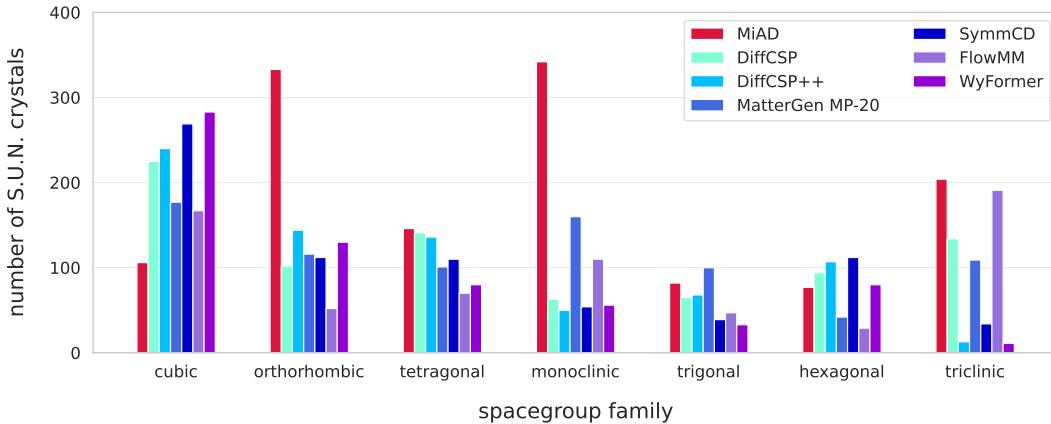


Figure 5: Comparison of models for de novo crystal generation in terms of numbers of S.U.N. crystals belonging to main spacegroup families. We categorize only S.U.N. crystals produced by generative models with a fixed budget of 10 000 generated crystals. After generation, crystals are prerelaxed via CHGNet. Stability is also estimated via CHGNet. Spacegroup families are identified via SpacegroupAnalyzer from the pymatgen package with tolerance: 0.1.

Table 7: **Comparison of MiAD and MatterGen trained on Alex-MP20** We report stability and S.U.N. estimated eq-V2 for 10 000 sampled crystals. For clarity in evaluating the model’s quality, we also report the Unique&Novel rate among stable crystals.

Model	eq-V2 ($E^{\text{hull}} < 0.0$)		
	Stable (%) ↑	Unique&Novel (%) ↑	S.U.N. (%) ↑
MatterGen	5.2	39.4	2.1
MiAD	7.1	35.3	2.5

H COMPARISON ON ALEX-MP20

To demonstrate the scalability of MiAD, we conduct experiments on the large-scale Alex-MP20 dataset (Zeni et al., 2024) and compare its performance with that of MatterGen. We employ S.U.N., estimate stability using eq-V2, and use the same energy hull as in the previous experiments. The only difference in the evaluation protocol is that novelty is measured relative to the Alex-MP20 training set. Samples from both models are pre-relaxed for 100 steps using eq-V2. MiAD is trained for 1200 epochs with the same hyperparameters as in the MP-20 experiments. The results in Table 7 demonstrate that MiAD, without additional hyperparameter tuning, achieves state-of-the-art performance on one of the largest datasets for de novo crystal generation.

I LLM USAGE

The text of the paper was polished for grammar and style using LLMs.

J QUALITATIVE ANALYSIS OF MIRAGE INFUSION DYNAMICS

The proposed MiAD framework generalizes standard crystal diffusion models by introducing dynamic atom counts. This approach differs from fixed-size baselines along three primary axes: (1) the formulation of the generative task, (2) the robustness of generation trajectories, and (3) the expanded action space available during diffusion. While the first two axes are implicit, the third, the ability to dynamically add or remove atoms, enables a direct analysis of how the model manages structural evolution. In this section, we investigate the hypothesis that mirage atoms provide an “error correction” mechanism, allowing the model to recover from unstable intermediate states.

1188

1189

PreRelaxed(M_t) : $\text{Yb}_3\text{BAsBr}_3$

1190

1191

1192

1193

1194

1195

1196

1197

1198

1199

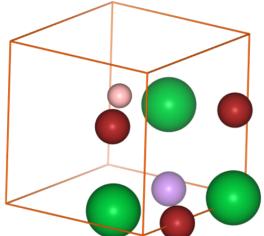
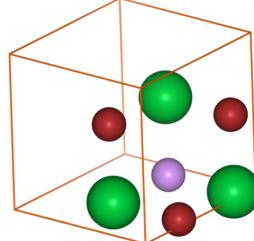
Mirage was added at $t = 95$ PreRelaxed(\hat{M}_t) : Yb_3AsBr_3 

Figure 6: Comparison between the counterfactual structure M_t (left), and the corrected structure \hat{M}_t (right). The corrected structure is thermodynamically stable with a higher symmetry group (monoclinic) compared to the unstable, triclinic counterfactual.

1200

1201

1202

1203

1204

1205

J.1 METHODOLOGY

1206

1207

To isolate the impact of the mirage mechanism, we focused on "boundary" cases in the generation trajectories where the model converts a real atom into a mirage atom (effectively removing it) during the final stages of the reverse diffusion process. We performed a comparative analysis on these samples, defined as follows:

1211

1212

The Counterfactual (M_t): The structure at the moment right before the last change in the denoising trajectory from a real atom to a mirage atom. In this state, the atom remains a real element, simulating a fixed-size model that cannot delete atoms, even when they are structurally disadvantageous.

1213

1214

The Corrected Structure (\hat{M}_t): The crystal structure with the same set of atoms as right after the model successfully changed the real atom to the mirage atom. To ensure that differences arise solely from the composition change (and not from stochastic noise or drift), we preserve the exact lattice parameters and fractional coordinates of all remaining atoms from M_t .

1215

1216

1217

1218

Both M_t and \hat{M}_t were subsequently pre-relaxed using CHGNet to evaluate their stability and symmetry properties.

1219

1220

1221

1222

J.2 EVIDENCE OF ERROR CORRECTION

1223

1224

Our observations reveal that the mirage mechanism functions as a dynamic pruning tool. Out of 1000 generated crystals, we analyzed a subset of 57 that were identified as Stable, Unique, and Novel, where the transition from a real atom to a mirage occurred after 150 steps of the denoising process (out of 1000).

1225

1226

The removal of the superfluous atoms in these trajectories resulted in drastic improvements in thermodynamic stability:

1227

1228

1229

1230

1231

1232

1233

1234

1235

1236

1237

1238

1239

1240

1241

1. The median energy above hull (E^{hull}) dropped from 0.29 eV/atom for the counterfactuals (M_t) to 0.033 eV/atom for the corrected structures (\hat{M}_t).
2. The median absolute deviation of the energy narrowed significantly from 0.27 to 0.12, indicating a more consistent convergence toward stable minima.

Furthermore, the structural symmetry improved substantially. Among the 57 crystals, only 14 exhibited non-trivial symmetry groups (cubic, monoclinic, orthorhombic, triclinic, or trigonal) in the uncorrected state. After the mirage infusion, this number nearly doubled to 27 crystals. This suggests that the model effectively removes atoms that break symmetry or disrupt the lattice, a capability structurally impossible for fixed-size diffusion baselines.

1242
1243
1244
1245
1246Table 8: **Scaling mirage infusion to MPTS-52** MiAD is trained with $N_m = 57$. MiAD-WRNA is trained with $N_m = N_{\text{atoms}} + k$, where $k \sim \mathcal{U}[0, 10]$ (+5 mirage atoms in average). The models are compared using S.U.N. for 10 000 sampled crystals, where stability is estimated via (left) eq-V2 and (right) CHGNet.1247
1248
1249
1250
1251

Model	eq-V2 ($E^{\text{hull}} < 0.0$)			CHGNet ($E^{\text{hull}} < 0.0$)		
	Stable (%) ↑	Unique&Novel (%) ↑	S.U.N. (%) ↑	Stable (%) ↑	Unique&Novel (%) ↑	S.U.N. (%) ↑
DiffCSP	4.3	37.2	1.6	11.8	69.9	8.2
MiAD	5.7	55.6	3.2	13.0	76.2	9.9
MiAD-WRNA	5.6	44.1	2.5	13.0	70.9	9.2

1252
1253
1254

J.3 CASE STUDY

1255
1256
1257
1258
1259
1260
1261

A concrete example illustrates this behavior: in one trajectory (see Figure 6), the model removed a Boron atom at step 95 (of 1000), shifting the composition from $\text{Yb}_3\text{BAsBr}_3$ (M_t) to Yb_3AsBr_3 (\hat{M}_t). While both structures initially possessed triclinic symmetry, the pre-relaxed crystal \hat{M}_t converged to a higher-symmetry monoclinic group and achieved thermodynamic stability (negative energy above hull, -0.083 eV/atom). Conversely, the counterfactual M_t failed to find higher symmetry, remaining triclinic, and was found to be heavily unstable (0.335 eV/atom).

1262
1263

This confirms that the mirage atoms allow the model to alleviate mistakes made during earlier stages of generation, dynamically refining the composition to ensure realizability.

1264
1265
1266

K SCALABILITY ON CRYSTALS WITH LARGER NUMBER OF ATOMS

1267
1268
1269
1270
1271
1272

Mirage infusion enables a joint diffusion model operating on the (L, F, A) crystal representation to change the number of atoms in a crystal during generation. To assess how this procedure scales when the admissible range of atom counts is wider, we conduct experiments on MPTS-52 – a more challenging extension of MP-20 (Jain et al., 2013a), comprising 40,476 structures with up to 52 atoms per cell, ordered by earliest publication year in the literature.

1273
1274
1275
1276
1277
1278

There are several challenges in conducting comparisons on MPTS-52. To our knowledge, prior works have not reported wide comparisons of S.U.N. metrics for different generative models on MPTS-52. Metrics such as Structure Validity, Compositional Validity, Coverage (COV-R, COV-P), and distributional distances (e.g., the Wasserstein distance of scalar material properties such as density or number of elements) proposed in (Xie et al., 2022) have significant limitations, which we discuss in Appendix E.

1279
1280
1281
1282
1283

It is important to consider the DiffCSP backbone architecture, which operates most effectively with crystals containing fewer than 20 atoms. Applying DiffCSP to MPTS-52 places the model in a regime where performance degrades due to the larger number of atoms. Applying mirage infusion further increases the number of atoms; thus, when we use the same network architecture, that challenge is amplified.

1284
1285
1286
1287

We trained the original DiffCSP and MiAD on MPTS-52 for 6k epochs (six times the number of epochs used in the original DiffCSP for MP-20) and selected the best checkpoints by S.U.N. We use the same neural network as in experiments on MP-20. We did not tune MiAD for MPTS-52, but instead used the configuration identified as optimal on MP-20:

1288
1289

MP-20: $N_m = \text{maximum_number_of_atoms_in_mp20} + 5 = 20 + 5 = 25$

1290

MPTS-52: $N_m = \text{maximum_number_of_atoms_in_mpts52} + 5 = 52 + 5 = 57$

1291
1292
1293
1294
1295

The results are presented in Table 8. Here, we observe consistent improvements with mirage infusion using the same configuration (without finding optimal hyperparameters for dataset). Computational overhead: 1 epoch of DiffCSP training on MPTS-52: 35 sec, while 1 epoch of MiAD ($N_m = 57$) training on MPTS-52: 79 sec ($\approx 2.26x$ compared to DiffCSP). The overhead is smaller than in MP-20 (see Appendix C) because the network is already near its upper limit of effective atomic interactions (20 atoms); further increases in atom count then lead to approximately linear, not quadratic, scaling.

1296 Applying mirage infusion to datasets with a broad range of atom counts is a plausible use case in
1297 practical applications. For such datasets, we can vary the number of mirage atoms added per structure
1298 while keeping the rest of the procedure unchanged. We demonstrate this variant on MPTS-52.
1299 Specifically, we set N_m per crystal to $N_{\text{atoms}} + \text{Uniform}[0, 10]$, i.e., each crystal receives 0-10 mirage
1300 atoms (+5 on average). At generation start, we sample the number of atoms from the training-set
1301 distribution (as in the original DiffCSP), and add a sample from $\text{Uniform}[0, 10]$. The results for
1302 this variant (MiAD-WRNA: MiAD for Wide Ranges of Number of Atoms) presented in Table 8.
1303 MiAD-WRNA still provides substantial gains over DiffCSP, although it performs slightly worse than
1304 the original MiAD with fixed N_m . Its main advantage lies in computational cost: 1 epoch of DiffCSP
1305 training on MPTS-52: 35 sec; 1 epoch of MiAD ($N_m = 57$) training on MPTS-52: 79 sec ($\approx 2.26x$
1306 vs DiffCSP); 1 epoch of MiAD-WRNA (+5 avg) training on MPTS-52: 40 sec ($\approx 1.14x$ vs DiffCSP).
1307 Thus, MiAD-WRNA offers a practical trade-off between computational cost and quality.

1308 These MPTS-52 experiments, together with Table 3, support the effectiveness of mirage atoms across
1309 different strategies for incorporating them into crystals.

1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349