



**HAL**  
open science

# Vector Field Oriented Diffusion Model for Crystal Material Generation

Astrid Klipfel, Yaël Fregier, Adlane Sayede, Zied Bouraoui

► **To cite this version:**

Astrid Klipfel, Yaël Fregier, Adlane Sayede, Zied Bouraoui. Vector Field Oriented Diffusion Model for Crystal Material Generation. The 38th Annual AAAI Conference on Artificial Intelligence, Feb 2024, Vancouver, Canada. pp.22193-22201, 10.1609/aaai.v38i20.30224 . hal-04526745

**HAL Id: hal-04526745**

**<https://univ-artois.hal.science/hal-04526745>**

Submitted on 29 Mar 2024

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Vector Field Oriented Diffusion Model for Crystal Material Generation

Astrid Klipfel<sup>1,2,3</sup>, Yaël Fregier<sup>2</sup>, Adlane Sayede<sup>3</sup>, Zied Bouraoui<sup>1</sup>

<sup>1</sup>Univ. Artois, UMR 8188, Centre de Recherche en Informatique de Lens (CRIL), F-62300 Lens, France.

<sup>2</sup>Univ. Artois, UR 2462, Laboratoire de Mathématiques de Lens (LML), F-62300 Lens, France.

<sup>3</sup>Univ. Artois, UMR 8181, Unité de Catalyse et de Chimie du Solide (UCCS), F-62300 Lens, France.  
{astrid.klipfel, yael.fregier, adlane.sayede, zied.bouraoui}@univ-artois.fr

## Abstract

Discovering crystal structures with specific chemical properties has become an increasingly important focus in material science. However, current models are limited in their ability to generate new crystal lattices, as they only consider atomic positions or chemical composition. To address this issue, we propose a probabilistic diffusion model that utilizes a geometrically equivariant GNN to consider atomic positions and crystal lattices jointly. To evaluate the effectiveness of our model, we introduce a new generation metric inspired by Frechet Inception Distance, but based on GNN energy prediction rather than InceptionV3 used in computer vision. In addition to commonly used metrics like validity, which assesses the plausibility of a structure, this new metric offers a more comprehensive evaluation of our model’s capabilities. Our experiments on existing benchmarks show the significance of our diffusion model. We also show that our method can effectively learn meaningful representations.

## Introduction

Crystal materials play a crucial role in numerous technological applications, such as the aerospace industry and semiconductors. As global warming poses a significant threat to our society, developing new materials can be vital to our ability to adapt to change and minimize our carbon footprint. To support the ecological transition, it is important to search for and generate new crystals, such as new semiconductors, that can be used to produce solar hydrogen or store hydrogen in solid form (Tournet et al. 2020; Rusman and Dahari 2016). New methods, such as high-throughput screening, rely heavily on machine learning models to speed up the development process. Recently, there has been growing interest in developing machine learning models to create new crystals (Noh et al. 2019; Dan et al. 2020; Long et al. 2021). In particular, Graph-based Neural Networks (GNNs) models have been proven effective for handling crystal structures. While, however, these models have shown excellent performances for regression tasks and property prediction such as the formation energy or the band-gap (Xie and Grossman 2018; Jørgensen, Jacobsen, and Schmidt 2018; Choudhary and DeCost 2021), the progress is still hampered by the limited extent to which generation is performed. As far as we

know, no existing model can perform sampling with a precise composition. This type of sampling is crucial for exploring the various phases of a material within a specific composition and generating a convex hull. The ability to do this is essential in searching for new materials (Kim et al. 2020). Our work is unique in enabling this type of sampling, whereas other models are not as well-suited for this task. In this paper, we propose an effective probabilistic sampling diffusion process with GNN for crystal materials generation.

Crystals are periodic structures consisting of a minimal set of atoms (unit cell) infinitely repeated in all directions of 3D space. Contrary to organic chemistry, where diffusion models have led to spectacular results on generation, a key challenge when dealing with crystals is how the periodicity and the complex chemistry of materials are taken into account and handled by the generative model. In (Xie et al. 2021), a GNN-based model for material generation has been proposed where the diffusion process is only applied to atomic positions but not cells. Generating an appropriate cell is crucial for crystals. A cell defines the repetition pattern of a material and its density, which have an important impact on the properties of the crystal lattice. To this end, we propose a diffusion process on all the geometry of the structures (atomic positions and cells). In particular, our model relies on a diffusion process in a torus (for the atoms) and on a lattice (for the cells). We use a neural network based on an equivariant vector field to create our lattices, which sets us apart from others who only predict lattice parameters from geometry. Additionally, we use an equivariant graph neural network layer. This layer is equivariant to both  $\mathbb{E}$  and  $SL_3(\mathbb{Z})$ . Although introducing this diffusion process is challenging, we offer an elegant solution and explain how to reverse the process. Our model generates a more realistic lattice, as confirmed by our experiments.

To assess the quality of the generated structure, several metrics such as validity, density and formation energy are commonly used (Xie et al. 2021; Court et al. 2020). The validity metric performs well when the generated crystals are bad, i.e., to discard invalid materials having a poor validity ratio. When a generative model produces a good crystal structure, this metric is no longer meaningful. The density and formation energy are good metrics. However, they are arbitrary and partially cover material characteristics. There exist also metrics that compare the distributions of other

physical quantities. Contrary to the validity metric, the comparison of statistical distributions is a good metric. However, even though the compared quantities are carefully chosen such as density and energy, they are not enough to assess the quality of generated materials. Namely, other quantities such as band-gap, magnetism, piezoelectricity and elasticity are also relevant for characterizing materials. To this end, in addition to existing metrics, we add Frechet ALIGNN Distance (FAD), which provides a more complete view for evaluating crystals as it takes into account a large number of features. This metric aims to provide a robust evaluation that is similar to the FID used for computer vision but using ALIGNN (Choudhary and DeCost 2021) instead of inception v3.

The main contributions of this paper are as follows: (i) we propose a diffusion process on the whole geometry of crystal materials; (ii) we introduce a regression model that learns the reverse process and generates materials; (iii) we propose a new metric to evaluate the performance of material generative models; (iv) we performed deep experimental analysis of our models on existing benchmarks; (v) we provide the code to train our proposed model with diffusion and sampling of crystalline structures; (vi) we provide a tool to calculate the generation metrics as an independent python package <sup>1</sup>.

## Related Works

Recent years have witnessed an important interest in developing machine learning models for material science in both organic chemistry (Satorras, Hoogeboom, and Welling 2021; Igashov et al. 2022; Schneuing et al. 2022; Luo, Yan, and Ji 2021; Shi\* et al. 2020; Wu et al. 2021; Xu et al. 2022; Trippe et al. 2023) and materials science (Xie and Grossman 2018; Jørgensen, Jacobsen, and Schmidt 2018; Choudhary and DeCost 2021; Klipfel et al. 2023a). Most state-of-the-art generative methods can be divided according to the material representation: fingerprint (Kim et al. 2020; Ren et al. 2022; Noura, Sokolovska, and Crivello 2019; Dan et al. 2020; Zhao et al. 2021; Pathak et al. 2020; Sawada, Morikawa, and Fujii 2019; Hu et al. 2021), voxel (Long et al. 2021; Noh et al. 2019; Kim, Lee, and Kim 2020; Court et al. 2020; Hoffmann et al. 2019) and graph-based representation (Xie et al. 2021; Ekström Kelvinius, Armiento, and Lindsten 2022; Gibson, Hire, and Hennig 2022; Cheon et al. 2020). The fingerprint or voxel-based models are, in general, built upon VAE (Ren et al. 2022; Long et al. 2021; Noh et al. 2019) and GAN (Kim et al. 2020; Long et al. 2021; Zhao et al. 2021). Although these models perform well, generated structures often lack stability and are subject to different constraints. For example, the generated crystals are often limited to cubic cells (Zhao et al. 2021; Long et al. 2021; Chen and Pao 2019; Court et al. 2020), or have specific chemical composition (Kim et al. 2020; Long et al. 2021; Noh et al. 2019). With such constraints, a lot of data is required on specific chemical composition and learned knowledge cannot be easily transferred and reused. Moreover, Voxel-based models are, in general, very expensive to train and difficult to use on large volumes of data. Finally,

<sup>1</sup>Supplementary materials are available at [arxiv.org/abs/2401.05402](https://arxiv.org/abs/2401.05402), source code are available at: [github.com/aklipfel/gemsdiff](https://github.com/aklipfel/gemsdiff)

with a diffusion process, we can often have very good performance in generation while avoiding the risk of collapse mode (the main problem with GAN-based models).

Recently, graph-based models with diffusion process have shown impressive results in generation for organic chemistry (Schneuing et al. 2022; Igashov et al. 2022; Wu et al. 2021). However, there are few works for crystal generation (Xie et al. 2021). There are only a few models that can accurately predict lattice while remaining invariant to the  $\mathbb{E}$  group (Xie et al. 2021; Yan et al. 2022). However, these models tend to not be invariant to the  $SL_3(\mathbb{Z})$  group, which may limit their prediction capabilities. In (Xie et al. 2021), the diffusion process does not apply to cells and only concerns atomic positions and atomic numbers. It is not designed to sample a given chemical composition, i.e., sampling structures for various atomic ratios. This is a clear limitation when constructing the convex hull of a given chemical composition. Moreover, the training phase suffers from instability since the diffusion is not clearly defined in the torus or on the cell, which consequently lowers the precision of the predicted lattice parameters. These limitations prevent the application of such diffusion models on real-world problems, i.e., when studying chemical systems for a specific application. Our goal is to create a model that can generate structures with specific chemical compositions but with varying proportions. To accomplish this, our model allows us to perform sampling by adjusting the rate given as input.

## Background

We first introduce background elements about crystalline materials and the probabilistic diffusion process.

### Crystalline Materials

Crystals are structures made up of a minimum number of atoms called unit cell that are repeated infinitely in all directions of space to form a lattice. This lattice can be seen as an atomic point cloud, where an atom may be repeated in multiple positions due to the space tiling. Hence, its local environment can overlap with adjacent repetitions. To define a crystal, we use the atomic positions  $x_i$ , which are within the range of  $[0, 1]^3$ , the chemical feature space  $F$  that represents the chemical information of each atom  $z_i \in F$ , and the lattice  $\rho \in GL_3(\mathbb{R})$  that represents the periodicity. The infinite point cloud can be expressed as follows:

$$\langle M \rangle = \{(\rho(x_i + \tau), z_i) | \tau \in \mathbb{Z}^d, i \in [1..n]\} \subseteq \mathbb{R}^d \times F \quad (1)$$

Where  $\tau$  acts as a  $\mathbb{Z}^3$  vector that translates the point cloud. The representation space of all possible materials with  $n$  atoms is then defined as  $\mathcal{M}^F = \bigcup_{n \in \mathbb{N}} \mathcal{M}_n^F$  with

$$\mathcal{M}_n^F = \{(\rho, x, z) | \rho \in GL_3(\mathbb{R}), x \in [0, 1]^{n \times d}, z \in F^n\} \quad (2)$$

The cloud  $\langle M \rangle$  can be rotated and translated by the Euclidian group. However, we can also find an alternative tiling of the space. Let  $\langle M \rangle$  invariant in space, which corresponds to an action of  $SL_3(\mathbb{Z})$  on the materials, i.e., invariant under the action of the lattice  $L = \rho \cdot \mathbb{Z}^d \subseteq \mathbb{R}^d$ . We recall action

groups that a GNN model should satisfy when manipulating crystal lattice (Klipfel et al. 2023a): (i)  $O(d)$  orthogonal group acting by  $g \cdot (\rho, x, z) = (g \cdot \rho, x, z)$ ; (ii) the translation group  $\mathbb{R}^d$  acting by  $v \cdot (\rho, x, z) = (\rho, [x + v], z)$ ; (iii) the euclidean group  $E(d) = \mathbb{R}^d \rtimes O(d)$ ; (iv) the special linear group  $SL_d(\mathbb{Z})$  acting by  $g' \cdot (\rho, x, z) = (\rho g'^{-1}, [g'x], z)$ . We refer to (Klipfel et al. 2023a) for more details.

## Probabilistic Diffusion

Diffusion models are generative models developed according to the nonequilibrium thermodynamics (Sohl-Dickstein et al. 2015; Ho, Jain, and Abbeel 2020), which define Markovian diffusion process that incrementally introduces noises to input data, and then learns from a reverse process to distinguish data from noises. While it is relatively easy to add noises to input data, such an operation generally engenders information loss since several identical antecedents to the function that added noises may hold. The *forward process* or *diffusion process* is defined as a Markov chain that incrementally adds noises to the  $x_0$  data. The diffusion process is the *posterior* of a latent variable model and is usually defined using a scheduler of variance  $\beta_1, \dots, \beta_T$ :

$$\begin{aligned} q(x_t|x_{t-1}) &:= \mathcal{N}(\sqrt{1 - \beta_t}x_{t-1}, \beta_t \mathbf{I}), \\ q(x_{1:T}|x_0) &= \prod_{t=1}^T q(x_t|x_{t-1}) \end{aligned} \quad (3)$$

The *reverse process* and the *joint distribution* of the latent variable with  $p(x_T) = \mathcal{N}(0, \mathbf{I})$  can similarly be defined as a Markov chain:

$$\begin{aligned} p_\theta(x_{t-1}|x_t) &:= \mathcal{N}(\mu_\theta(x_t, t), \Sigma_\theta(x_t, t)), \\ p_\theta(x_{0:T}) &= p(x_T) \prod_{t=1}^T p_\theta(x_{t-1}|x_t) \end{aligned} \quad (4)$$

Similar to variational models, diffusion models can be trained by minimizing the evidence lower bound (ELBO):

$$\begin{aligned} \mathcal{L} &= \mathbb{E}_q \left[ \underbrace{D_{\text{KL}}(q(x_T|x_0)||p_\theta(x_T))}_{L_T} \right. \\ &+ \left. \sum_{t>1} \underbrace{D_{\text{KL}}(q(x_{t-1}|x_t, x_0)||p_\theta(x_{t-1}|x_t))}_{L_{t-1}} - \underbrace{\log p_\theta(x_0|x_1)}_{L_0} \right] \end{aligned} \quad (5)$$

Equation 5 can be simplified when the Markov chain is defined with Gaussian distributions and the variance scheduler only depends on hyperparameters.

$$\mathcal{L}_{\text{simple}} := \mathbb{E}_q \left[ \|\tilde{\mu}_t(x_t, x_0) - \mu_\theta(x_t, t)\|^2 \right] \quad (6)$$

with  $\tilde{\mu}_t(x_t, x_0) := \frac{\sqrt{\bar{\alpha}_t - 1} \beta_t}{1 - \bar{\alpha}} x_0 + \frac{\sqrt{\bar{\alpha}_t(1 - \bar{\alpha}_{t-1})}}{1 - \bar{\alpha}_t} x_t$ ,  $\alpha_t := 1 - \beta_t$ ,  $\bar{\alpha}_t := \prod_{s=1}^t \alpha_s$  and  $x_t(x_0, \epsilon) = \sqrt{\bar{\alpha}_t} x_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon$  with  $\epsilon \sim \mathcal{N}(0, \mathbf{I})$ .

## Diffusion on a Torus with an Equivariant Graph Neural Network

We need two steps to apply diffusion to crystal generation. First, we define a diffusion on torus to handle the periodicity of a material. Second, we introduce a model  $f_\theta$  that

can learn the reverse process. We use an Equivariant Graph Neural Network  $f_\theta : \mathcal{M}_n^F \times \mathbb{N} \rightarrow \mathcal{M}_n^F$  that takes as input a crystal  $M_t = (\rho, x, z) \in \mathcal{M}_n^F$  at a diffusion step  $t \in \mathbb{N}$  and predicts a new cell  $\rho'$  as well as new atomic positions  $x'$ . This approximation of  $M_{t-1}$  enables to denoise the step  $t$  of the diffusion process. Since a crystal lattice forms a toroidal space, we adapt diffusion to such a space using a loss of two parts: A first part concerns the shape of the torus  $\mathcal{L}^\rho$ , and the second part concerns the diffusion inside the torus  $\mathcal{L}^x$ . The final loss is the sum of the two parts:  $\mathcal{L} := \mathcal{L}^\rho + \mathcal{L}^x$ . We now describe these different steps in more detail.

## Equivariant Graph Neural Network

We start by introducing the map  $f_\theta$  which is an equivariant GNN that takes a material and a diffusion step as input and outputs a new material with the same chemical composition but an updated lattice  $\rho$  and atomic position  $x$ .

$$\begin{aligned} f_\theta : \mathcal{M}_n^{\mathbb{A}} \times \mathbb{N} &\rightarrow \mathcal{M}_n^{\mathbb{A}} \\ (\rho, x, z, t) &\mapsto (\rho', x', z) \end{aligned}$$

The goal of the map  $f_\theta$  is to enable inverting the diffusion process one step at a time by parameterizing the  $p_\theta(M_{t-1}|M_t)$ . This map is then defined with a step-dependent auto-encoder GNN.

$$\begin{aligned} \text{GNN}_{AE} : \text{GL}_3(\mathbb{R}) \times \mathbb{T}^n \times F^n \times \mathbb{N} &\rightarrow \text{GL}_3(\mathbb{R}) \times \mathbb{R}^{n \times 3} \times F^n \\ (\rho, x, z, t) &\mapsto (y^\rho, y^x, y^z) \end{aligned}$$

The pair of action  $(y^\rho, y^x)$  predicted by the GNN can be used to update the geometry of the material as follows:

$$(\rho', x', z) = (y^\rho, y^x, \text{id}) \cdot (\rho, x, z) \text{ with } \begin{cases} \rho' := y^\rho \rho \\ x'_i := [x_i + y_i^x]. \end{cases} \quad (7)$$

The auto-encoder defining the denoising task consists of the composition of three maps:

$$\text{GNN}_{AE} := \text{Decod} \circ \text{Encod} \circ \text{emb}.$$

The first map,  $\text{emb}$ , embeds the atoms for the network. This embedding is built from the atomic number and the diffusion step as  $\text{emb} : \mathbb{A} \times \mathbb{N} \rightarrow F$ , given by the following two cases:

$$\begin{cases} \text{emb}(z, t)_{2k} &= \sin\left(\frac{t}{n^{2k/d_{\text{emb}}}}\right) + \text{emb}_{z, 2k} \\ \text{emb}(z, t)_{2k+1} &= \cos\left(\frac{t}{n^{2k/d_{\text{emb}}}}\right) + \text{emb}_{z, 2k+1} \end{cases} \quad (8)$$

With  $\text{emb}_z$  is a vector associated with the atomic number  $z$ . The embedding works by associating a vector to each atomic number and adding  $t$ , the index of the diffusion step, to this representation in the same way that positional encoding works in a transform on tokens. The remaining maps  $\text{Encod}$  and  $\text{Decod}$  are built with components of  $\text{GNN}_{\text{enc}}$  and  $\text{GNN}_{\text{dec}}$ , two equivariant GNN architectures of the type (Klipfel et al. 2023b). These components are combined as follows:

$$z_i^{\text{emb}} = \text{emb}(z_i, t) \quad (9a)$$

$$(\cdot, \cdot, z^{\text{enc}}) = \text{GNN}_{\text{enc}}(\rho, x, z^{\text{emb}}) \quad (9b)$$

$$(\rho', x', \cdot) = \text{GNN}_{\text{dec}}(\rho, x, z^{\text{enc}}) \cdot (\rho, x, \cdot) \quad (9c)$$

We show that  $f_\theta$  is equivariant (proof in Appendix A).

$$f_\theta(g \cdot M, t) = g \cdot f_\theta(M, t) \forall g \in E(3) \times \text{SL}_3(\mathbb{Z}). \quad (10)$$

In practice, we will use this architecture, but we decode  $\rho'$  from a cubic cell of volume 1 rather than from  $\rho$  because we found that more stable and produces better results. Notice that while this modification breaks the equivariance between the input and the output of the network, the GNN that constitutes it remains equivariant. As such, we still benefit from the GNN structuring factor. So,  $\rho'$  is computed as follows:

$$(\rho', x', \cdot) = \text{GNN}_{\text{dec}}(\rho, x, z^{\text{enc}}) \cdot (\mathbf{I}_3, x, \cdot). \quad (11)$$

## Diffusion Process for Crystals

We first propose a diffusion on  $\text{GL}_3(\mathbb{R})$ , the space of lattices. As the space is non-linear, we linearize with logarithm to use the diffusion process of  $\mathbb{R}^n$  and bring back the result of the diffusion to  $\text{GL}_3(\mathbb{R})$  using the exponential map. We then show that the usual diffusion process in the Euclidian space induces a diffusion process on the torus, enabling us to also apply diffusion to atoms within a lattice.

**Lattice Vector Space.** One can see a cell  $\rho \in \text{GL}_3(\mathbb{R})$  as the image by the exponential map of an element  $g \in \mathfrak{gl}_3(\mathbb{R})$ , i.e. as  $\rho = \exp(g)$ . By choosing a supplement  $\mathfrak{sl}_3(\mathbb{R}) \setminus \mathfrak{so}_3(\mathbb{R}) \subset \mathfrak{sl}_3(\mathbb{R})$ , one can decompose the Lie algebra  $\mathfrak{gl}_3(\mathbb{R})$  as a direct sum:

$$\mathfrak{gl}_3(\mathbb{R}) = \mathfrak{so}_3(\mathbb{R}) \oplus \mathfrak{sl}_3(\mathbb{R}) \setminus \mathfrak{so}_3(\mathbb{R}) \oplus \mathbf{1}_3, \quad (12)$$

with

- $\mathfrak{so}_3(\mathbb{R}) \subset \mathfrak{gl}_3(\mathbb{R})$ : rotation algebra.
- $\mathfrak{sl}_3(\mathbb{R}) \setminus \mathfrak{so}_3(\mathbb{R}) \subset \mathfrak{gl}_3(\mathbb{R})$ : a subspace of the special linear algebra supplementary to  $\mathfrak{so}_3(\mathbb{R})$ . One can choose its elements to be responsible for the shape of the cell but with a cell of volume 1 ( $\det 1$ ).
- $\mathbf{1}_3 \subset \mathfrak{gl}_3(\mathbb{R})$ : the algebra composed of the identity matrix singleton, which generates the volume of the cell.

This decomposition of the linear Lie algebra allows us to define a lattice vector space where vectors are in 9 dimensions with 3 dimensions controlling the orientation of the cells, 5 dimensions for the shape of the cells and 1 dimension for volume. Since linear spaces are of the same finite dimension, they are isomorphic. A linear isomorphism is given by choosing a basis in  $\mathfrak{gl}_3(\mathbb{R})$ . An explicit basis  $\pi \in \mathfrak{gl}_3(\mathbb{R})^9$  defining the following isomorphism is given in Section B of the Appendix. Notice that this process is invertible.

$$\begin{aligned} \pi : \mathbb{R}^9 &\rightarrow \mathfrak{gl}_3(\mathbb{R}) \\ x &\mapsto \pi x. \end{aligned} \quad (13)$$

**Lattice Diffusion.** Based on Equation 6 and 13, we can define the loss in the lattice vector space where  $x_0^p$  is the vector of original lattice  $\rho_0$  and  $x_\theta(M_t, t)$  is the predicted lattice vector, resulting in:

$$\begin{aligned} \mathcal{L} &:= \mathbb{E}_q [ \|\tilde{\mu}_t(x_t^p, x_0^p) - \mu_\theta(M_t, t)\|^2 ] \\ &= \mathbb{E}_q \left[ \frac{\sqrt{\bar{\alpha}_{t-1}} \beta_t}{1 - \bar{\alpha}} \|x_0^p - x_\theta(M_t, t)\|^2 \right] \end{aligned} \quad (14)$$

Finally, we cannot directly use the loss defined in Equation 14 as the operations  $\exp$  and  $\log$  do not have a well-defined gradient in Pytorch due to the Eigenvalues decomposition. We, therefore, define a loss of lattice reconstruction that is invariant to the actions of  $\text{SO}_3(\mathbb{R})$ , which has the same minimum as the Equation 14 from the lattice predicted by the model  $\rho_\theta(M_t, t)$ .

$$\mathcal{L}^\rho := \mathbb{E}_q [ d(\exp(\pi \tilde{\mu}_t(x_t^p, x_0^p)), \rho_\theta(M_t, t)) ] \quad (15)$$

We choose as distance measure between the two lattices  $d(\rho, \rho') = \|p(\rho) - p(\rho')\|$  with  $p : \text{GL}_3(\mathbb{R}) \rightarrow \mathbb{R}^6$  as an application that associates the normalized crystallographic lattice parameters  $(a, b, c, \alpha, \beta, \gamma)$  to a lattice  $\rho$ .

**Torus Diffusion.** We now detail the forward process inside a torus and the loss function. The principle is similar to the classical diffusion process but with a periodic space.

*Forward Process Inside a Torus:* We define the torus as a quotient space such that  $\mathbb{T} = \mathbb{R}^d / \mathbb{Z}^d$  where we consider the embedding of  $\mathbb{Z}^d$  in  $\mathbb{R}^d$  given by integer coordinates. For a distribution  $P$  on  $\mathbb{R}^d$  of density  $p$ , we use  $[P]_{\mathbb{T}}$  to denote the distribution  $P$  in the torus with density  $p_{\mathbb{T}}(x) = \sum_{\tau \in \mathbb{Z}^d} p(x + \tau)$ . We show that if  $P$  is a density, then  $[P]_{\mathbb{T}}$  is a density (proof in Appendix C).

**Lemma 1**  $[\mathcal{N}(\mu, \sigma I)]_{\mathbb{T}} \xrightarrow{\sigma \rightarrow +\infty} \mathcal{U}(0, 1)^d$

The proof of Lemma 1 is provided in Appendix D. Consequently, we choose  $p_\theta(x_T) = \mathcal{U}(0, 1)$ . We can now introduce our forward process with  $q$  such that.

$$\begin{aligned} q(x_t | x_{t-1}) &:= [\mathcal{N}(x_0, \beta_t)]_{\mathbb{T}} \text{ then} \\ q(x_t | x_0) &= [\mathcal{N}(x_0, (1 - \bar{\alpha}_t) \mathbf{I})]_{\mathbb{T}} \end{aligned} \quad (16)$$

Notice that there is no need to shift the norm because the distribution tends towards a uniform distribution. Now the reverse process with  $p_\theta$  in a similar way to equation 4:

$$p_\theta(x_{t-1} | x_t) := [\mathcal{N}(\mu_\theta(x_t, t), \sigma_t^2 \mathbf{I})]_{\mathbb{T}} \quad (17)$$

We deduce that we can sample  $x_{t-1}$  from  $x_t$  and  $x_0$  directly as  $q(x_{t-1} | x_t, x_0) = [\mathcal{N}(x_0, \hat{\beta}_t \mathbf{I})]_{\mathbb{T}}$

*Loss Function:* The loss is composed of 3 parts from equation 5:  $L_T$ ,  $L_{t-1}$  and  $L_0$ .  $L_T$  is a constant because it depends only on  $\beta_t$  (variance scheduler).  $L_0$  is the same as  $L_{t-1}$  in our context. To perform the last diffusion step, we use  $L_{t-1}$  with a variance of 0. This leaves us  $L_{t-1}$ :

$$L_{t-1} = \mathbb{E}_q [ D_{\text{KL}}(q(x_{t-1} | x_t, x_0) \| p_\theta(x_{t-1} | x_t)) ] \quad (18)$$

Being periodic, the distributions  $q$  and  $p_\theta$  can shift by  $k \in \mathbb{Z}^3$ . We then obtain Equation 18 which is equivalent to the Equation 19. Proofs are given in Appendix E.

$$L_{t-1}^k := \mathbb{E}_{q,k} \left[ \frac{1}{2\sigma_t^2} \|x_0 - x_t - k - \epsilon_\theta(x_t, t)\|^2 \right] \quad (19)$$

We can now define the final loss. As optimizing for any  $k$  minimizes 18, we choose the  $k$  that minimizes the trajectory performed by the atom. We then attend the targeted position following the shortest path.

$$\begin{aligned} \mathcal{L}^x &:= \mathbb{E}_q [ \|x_0 - x_t - k^* - \epsilon_\theta(x_t, t)\|^2 ] \\ &\text{with } k^* = \arg \min_{k \in \mathbb{Z}^d} \|x_0 - x_t - k\| \end{aligned} \quad (20)$$

Finally, we obtain a reconstruction loss where the network is trained to move atoms back to their original position following the shortest possible path.

Parameter	Value
batch size	128
epochs	512
learning rate	1e-3
knn	32
diffusion steps	100
$(\beta_0^p, \beta_T^p)$	(1e-5, 1e-1)
$(\beta_0^x, \beta_T^x)$	(1e-6, 2e-3)

Table 1: Hyperparameters of GemsDiff

## Experiments

We present an evaluation of our proposed models GemsDiff. We will, in particular, focus on experiments about diffusion and evaluation based on Frechet distance. All experiments were conducted on a subset of Materials Project (Jain et al. 2013), which is a dataset that contains crystalline materials that have been studied with ab-initio calculations (physical simulation). We follow the same setting as (Xie et al. 2021) and filter the dataset by only keeping stable materials having less than 20 atoms. We only keep stable materials as we aim to assess the capability of our model to produce stable structures. We use the same training, validation and test splits as (Xie et al. 2021) for comparable results.

Notice that in our research, we have chosen not to evaluate Perov-5 and Carbon-24 for various reasons. Perov-5 is a dataset of perovskite structures with identical atomic positions but different compositions, which does not align with the purpose of our research, which is to predict geometry based on composition. Carbon-24 is a structured dataset simulated at 10GPa, which is equivalent to 100k atmospheric pressure, and prediction models from the literature cannot accurately predict the properties of materials under such different conditions. Lastly, our focus is on semiconductors, and we are not interested in carbon structures.

**Training** The different hyperparameters we used for training are reported in Table 1. The training is performed following Algorithm 1, which uses the GNN and loss functions provided in Section . The training time is estimated to 36 hours on an Nvidia RTX 8000 GPU with 48Go. The sampling of the test set, which contains 9046 crystals, takes about 45 minutes, which is normal for a diffusion model.

**Baselines.** We evaluate GemsDiff against the state-of-the-art baselines: CDVAE (Xie et al. 2021), FTCP (Ren et al. 2022), PGCGM (Zhao et al. 2023), G-SchNet (Gebauer, Gastegger, and Schütt 2019) and P-G-SchNet (Gebauer, Gastegger, and Schütt 2019). We re-trained the CDVAE and FTCP to sample new crystals and calculate the Frechet ALIGNN distance (FAD) (as described in section 5.2). For PGCGM (Zhao et al. 2023), we used a pre-trained model that allows the generation of crystal materials. For G-SchNet and P-G-SchNet, we report results from (Xie et al. 2021).

**Sampling Process.** To generate structures with GemsDiff, we use chemical compositions of the test set. As such, geometrical information and chemical compositions have not been seen by the models. We use the Langevin dynamics with Algorithm 2 following the same parameters as in the training. We generate a dataset having the same size as the

---

### Algorithm 1: Training

---

**Input:**  $(\rho_0, x_0, z) \in \mathcal{M}$

- 1: **repeat**
- 2:  $t \sim \mathcal{U}(\{1, \dots, T\})$
- 3:  $x_t^p \sim q(x_t^p | \pi^{-1} \log \rho_0)$
- 4:  $x_t \sim q(x_t | x_0)$
- 5:  $(\rho_\theta, \epsilon_\theta) \leftarrow f_\theta(\exp \pi x_t^p, x_t, z, t)$
- 6: gradient descent step on  $\nabla \mathcal{L}$
- 7: **until** converged

---



---

### Algorithm 2: Sampling

---

**Input:**  $z \in \mathbb{A}^n$

- 1:  $x_T \sim \mathcal{U}(0, 1)$
- 2:  $\rho_T \leftarrow \exp(\pi x_T^p), x_T^p \sim \mathcal{N}(0, 1)$
- 3: **for**  $t \leftarrow T, \dots, 1$  **do**
- 4:  $(x', \rho') \leftarrow f_\theta(\rho_t, x_t, z, t)$
- 5:  $x_{t-1} \leftarrow [x' + \sigma_t^x \epsilon], \epsilon \sim \mathcal{N}(0, 1)$
- 6:  $\mu_\rho \leftarrow \tilde{\mu}_t(\pi^{-1} \log \rho_t, \pi^{-1} \log \rho')$
- 7:  $\rho_{t-1} \leftarrow \exp(\pi(\mu_\rho + \sigma_t^p \epsilon)), \epsilon \sim \mathcal{N}(0, 1)$
- 8: **end for**

---

test set to be able to compare them. The data generation is performed from the checkpoint, which has the best loss in validation. We use an exponential moving average of the weights with a decay of 0.995. Then, we use a set of crystals generated without filtering and without optimizing the geometry (without simulation methods to improve the geometry). We also do not search the space group or symmetries of the structures to improve geometry.

We can see structures generated by GemsDiff and from Materials Project in Figures 1 and 2. We can observe several well-known lattice systems having different constraints: cubic cell, monoclinic, orthorhombic, tetragonal, and, more importantly, hexagonal with 60° and 120° angles as lattice parameters. It is important to notice that there is no explicit rule in our model that facilitates the generation of such constraints in the lattice parameters. Interestingly enough, many lattices fall in a known system with a precision of less than 0.5 Angstrom and 0.1 Angstrom. Moreover, we can see that atoms of the materials make patterns like alignment, cubic geometry, or circular ring composed of 6 atoms. One can observe that the geometrical shape are coherent, i.e. often composed of the same type of elements. We stress the fact that there are no explicit constraints encoded in our model to encourage obtaining such structures. It is an emerging behavior that appears during training. Moreover, the observed characteristic of real materials is present in the generated data without any explicit bias encoded in our model. This offers an excellent indicator of the quality of the generated structures.

## Evaluation on Validity, Earth Mover Distance (EMD) for Density and Energy

To assess the quality of the generated crystals, we first rely on existing metrics: Validity and Earth Mover distance (EMD) for density and energy. EMD is calculated from the structure generated without filtering where the same amount of structures as in the test set (9048) is generated. Validity

Method	Validity (%)	density (EMD $g/cm^3$ )	energy (EMD $eV/atom$ )
FTCP*	88.18	1.586	0.8281
PGCGM	99.03	0.89	1.51
G-SchNet	99.65	3.034	42.09
P-G-SchNet	77.51	4.04	2.448
CDVAE	<b>100.0</b>	0.6875	0.2778
GemsDiff	99.05±0.20	<b>0.3331±0.0327</b>	<b>0.1984±0.0086</b>

Table 2: Results for different models in terms of validity, density and energy. For FTCP baseline, we report the results that we obtained, which are better than the original value reported in (Xie et al. 2021).

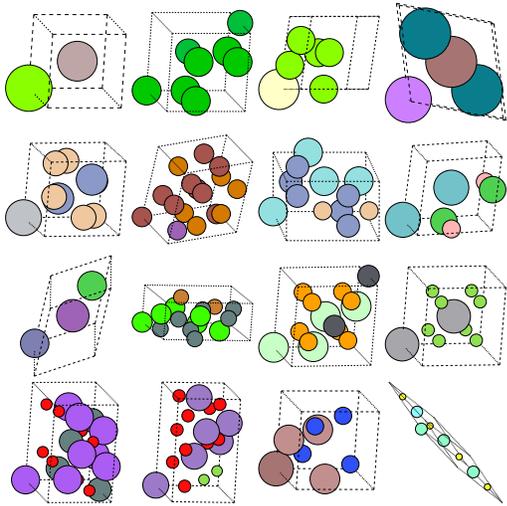


Figure 1: Sampling with GemsDiff and Langevin Dynamics

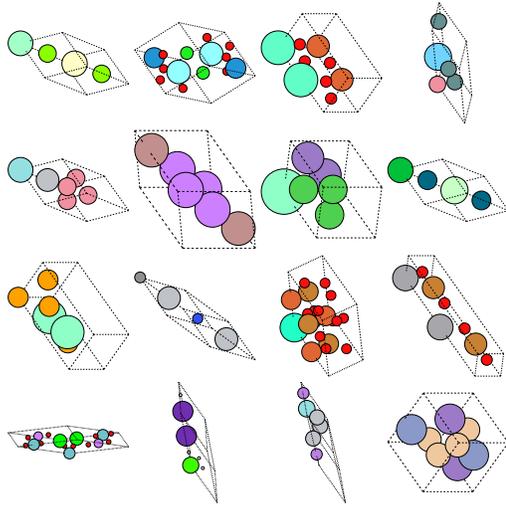


Figure 2: Crystals from materials project

(Court et al. 2020) refers to the percentage of structures that don't have two atoms with a relative distance less than 0.5 Angstrom. This metric evaluates the percentage of structures that are incoherent as physically having a distance lower than 0.5 Angstrom is impossible. Even if it is not a good metric, it at least allows us to spot the badly generated crystals. EMD, also called Wasserstein distance (or Wasserstein 1), allows the comparison of the distance between two statistical distributions. A generative model should be able to generate a structure distribution with properties that are close to its training distribution. So, the farther the distance between the distribution of properties of the generated structures to the distribution of the trained properties, the more different the generated data from the trained data. Consequently, the smaller the EMD, the better the result. We consider density distribution and energy distribution for the EMD metrics. As it is difficult to obtain and very expensive to simulate, some GNNs allow us to estimate it. We can use the same GNN for the Fréchet distance (ALIGNN) to estimate the energy (Section 5.2). The results are reported in Table 2. Our model outperforms the baselines on all the properties. EMD of the formation energy is significantly better than CDVAE. We also notice that the biggest gain is made for the density of the generated structures. We can explain this by the fact that our model does diffusion on the cell as well as on the atomic

positions. In this case, our model generates a volume that is more adapted to the composition of the structures. The validity is lower for our model than for CDVAE, but this metric is limited and our model remains very competitive.

### Fréchet Distance with ALIGNN (FAD)

We generate a vector of the latent space with a GNN (ALIGNN) to which we remove the last prediction layer and we compute the Fréchet distance as in (Heusel et al. 2017). We use a pre-trained model. Table 3 contains all FAD for the generated data. To test the relevance of the Fréchet distance, previous works on computer vision test the generated data with distortion or with different levels of scattering. In the case of materials, distortion and scattering are very similar. We perform a test for diffusion because it will also allow us to make a link between the FAD and the distance of the generated structures with equilibrium positions. To this end, we apply several levels of diffusion to see the associated FAD. We recall that the lower the FAD, the more similar the distributions are. So the lower the FAD, the more the set of tested materials contains materials close to their equilibrium. As we can see in Figure 3, the FAD is almost zero for  $t = 0$ , i.e. without diffusion. In contrast, the more  $t$  increases, the more the FAD increases. From a certain level of scattering, the FAD reaches a plateau and does not progress anymore. This

Method	Fréchet Distance
PGCGM	30.692
FTCP	16.562
CDVAE	1.777
GemsDiff	<b>1.509±0.115</b>

Table 3: To ensure consistent crystal quality in our training process, we repeat the training eight times and calculate the mean and standard deviation. This computation is carried out on 9048 structures without any filtering.

result is interesting because we observe a behavior close to the FID in the diffusion models used in computer vision (Ho, Jain, and Abbeel 2020). We also notice that the FAD reaches a lower plateau, i.e. the values are globally lower than for the FID in computer vision. This difference can be simply explained by the fact that the vectors used in inceptions are of size 2048 while the latent vectors produced by ALIGNN are only of size 256. To validate the number of instances between two samples, we compute the FAD by progressively increasing the sample size. We evaluate how the metric behaves and try to see when the FAD is close enough to 0. This is a trick that has already been done for the FID, and it allows us to be sure that the metric will not be too disturbed by statistical fluctuations. We can see the results for several sample sizes in figure 4. As for the FID, we can see that about 10000 samples are largely sufficient to have a reliable FAD measurement. So we will compare with sample sizes of 9046 (test set). In Table 3, we can see that our model has the best result for the Fréchet distance. This result consolidates the results of the other table with a new metric. To allow a better justification of our architecture, an ablation test is proposed in Appendix F.

## Discussion

Our work has the focus of prioritizing specialized tasks instead of generating compositions. Unlike other models that aim to create stable structures without composition constraints, our model is less suited for this purpose. However, chemists require models that can generate structures with certain compositions as they seek to find different metastable phases for a specific composition. Our model is specialized in generating compositions and, therefore, is more suitable for this type of task. Additionally, chemists are interested in determining the convex hulls of certain chemical compositions, which our model is very suitable for. Constraints on chemical composition, such as excluding rare and expensive elements, are also important considerations for chemists. Although some models are good at generating chemical compositions, it can be positive not to be able to generate compositions in certain cases. By breaking down the lattice in a Lie algebra and producing the lattice from an invariant vector field, the generation of the lattice appears to be enhanced. As Table 3 shows, there is a significant improvement in density, making it much more comparable to the initial distribution. In the future, decomposing the lattice could have other applications. For instance, it may be feasi-

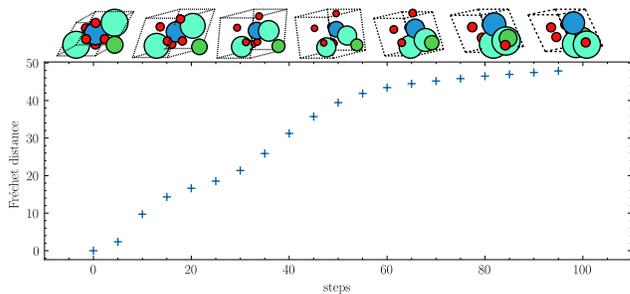


Figure 3: FAD and diffusion step.

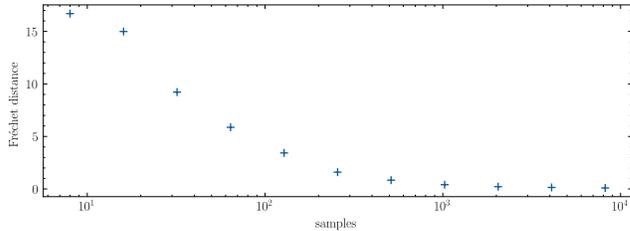


Figure 4: FAD for multiple sample sizes

ble to construct structures that are part of a particular space group and guarantee that the shape of the lattice is compatible with that space group using characteristics of the vector that represents the lattice in the Lie algebra’s space vector. The FAD aims to create a standardized metric for comparing crystal features, rather than relying on manually selected statistics. The features used are generated by the highly accurate ALIGNN, making them relevant for predicting material properties. However, assessing the consistency of the model’s generated features is challenging if ALIGNN is re-trained. To address this, we suggest using a pre-trainer version with consistent parameters and features for FAD calculation. To evaluate FAD quality, we can compare it to other metrics for generative model evaluation. Our experimental results show that FAD ranking aligns with other metrics, indicating its relevance for measuring generation quality.

## Conclusion

We have introduced a generative model for crystal materials with a diffusion process adapted to periodic structures based on an equivariant architecture that allows us to learn the reverse process. Our diffusion process performs on the whole crystal geometry (atomic position and cells) allowing us to improve the density prediction and hence obtain more realistic lattices compared to a simple model that predicts lattice parameters. Our diffusion model is particularly adapted to convex envelope sampling tasks, which is an important issue in materials science (Gubaev et al. 2019; Anelli et al. 2018). Finally, we have also proposed a new metric for evaluating crystal geometry inspired by works on computer vision which are based on a Fréchet distance and diffusion model. Our model produces interesting generated structures, each of which has a FAD that suggests its closeness to the equilibrium position (relatively stable structure).

## Acknowledgments

This work has been supported by ANR-22-CE23-0002 ERI-ANA, ANR-20-THIA-0004 and by HPC resources from GENCI-IDRIS (Grant 2022-[AD011013338] and 2021-[AD011013288]).

## References

- Anelli, A.; Engel, E. A.; Pickard, C. J.; and Ceriotti, M. 2018. Generalized convex hull construction for materials discovery. *Phys. Rev. Mater.*, 2: 103804.
- Chen, H.-A.; and Pao, C.-W. 2019. Fast and Accurate Artificial Neural Network Potential Model for MAPbI<sub>3</sub> Perovskite Materials. *ACS Omega*, 4(6): 10950–10959. PMID: 31460193.
- Cheon, G.; Yang, L.; McCloskey, K.; Reed, E. J.; and Cubuk, E. D. 2020. Crystal Structure Search with Random Relaxations Using Graph Networks. arXiv preprint arXiv:2012.02920.
- Choudhary, K.; and DeCost, B. 2021. Atomistic Line Graph Neural Network for improved materials property predictions. *npj Computational Materials*, 7(1): 185.
- Court, C. J.; Yildirim, B.; Jain, A.; and Cole, J. M. 2020. 3-D Inorganic Crystal Structure Generation and Property Prediction via Representation Learning. *Journal of Chemical Information and Modeling*, 60(10): 4518–4535. PMID: 32866381.
- Dan, Y.; Zhao, Y.; Li, X.; Li, S.; Hu, M.; and Hu, J. 2020. Generative adversarial networks (GAN) based efficient sampling of chemical composition space for inverse design of inorganic materials. *npj Computational Materials*, 6(1): 84.
- Ekström Kelvinius, F.; Armiento, R.; and Lindsten, F. 2022. Graph-based machine learning beyond stable materials and relaxed crystal structures. *Phys. Rev. Materials*, 6: 033801.
- Gebauer, N.; Gastegger, M.; and Schütt, K. 2019. Symmetry-adapted generation of 3d point sets for the targeted discovery of molecules. In Wallach, H.; Larochelle, H.; Beygelzimer, A.; d'Alché-Buc, F.; Fox, E.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Gibson, J.; Hire, A.; and Hennig, R. G. 2022. Data-augmentation for graph neural network learning of the relaxed energies of unrelaxed structures. *npj Computational Materials*, 8(1): 211.
- Gubaev, K.; Podryabinkin, E. V.; Hart, G. L.; and Shapeev, A. V. 2019. Accelerating high-throughput searches for new alloys with active learning of interatomic potentials. *Computational Materials Science*, 156: 148–156.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising Diffusion Probabilistic Models. In Larochelle, H.; Ranzato, M.; Hasselmin, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 6840–6851. Curran Associates, Inc.
- Hoffmann, J.; Maestrati, L.; Sawada, Y.; Tang, J.; Sellier, J. M. D.; and Bengio, Y. 2019. Data-Driven Approach to Encoding and Decoding 3-D Crystal Structures. *CoRR*, abs/1909.00949.
- Hu, J.; Yang, W.; Dong, R.; Li, Y.; Li, X.; Li, S.; and Sirlwardane, E. M. D. 2021. Contact map based crystal structure prediction using global optimization. *CrystEngComm*, 23: 1765–1776.
- Igashov, I.; Stärk, H.; Vignac, C.; Satorras, V. G.; Frossard, P.; Welling, M.; Bronstein, M.; and Correia, B. 2022. Equivariant 3D-Conditional Diffusion Models for Molecular Linker Design. arXiv:2210.05274.
- Jain, A.; Ong, S. P.; Hautier, G.; Chen, W.; Richards, W. D.; Dacek, S.; Cholia, S.; Gunter, D.; Skinner, D.; Ceder, G.; et al. 2013. Commentary: The Materials Project: A materials genome approach to accelerating materials innovation. *APL materials*, 1(1): 011002.
- Jørgensen, P. B.; Jacobsen, K. W.; and Schmidt, M. N. 2018. Neural Message Passing with Edge Updates for Predicting Properties of Molecules and Materials. arXiv:1806.03146.
- Kim, B.; Lee, S.; and Kim, J. 2020. Inverse design of porous materials using artificial neural networks. *Science Advances*, 6(1): eaax9324.
- Kim, S.; Noh, J.; Gu, G. H.; Aspuru-Guzik, A.; and Jung, Y. 2020. Generative Adversarial Networks for Crystal Structure Prediction. *ACS Central Science*, 6(8): 1412–1420. PMID: 32875082.
- Klipfel, A.; Bouraoui, Z.; Peltre, O.; Frégier, Y.; Harrati, N.; and Sayede, A. 2023a. Equivariant Message Passing Neural Network for Crystal Material Discovery. In Williams, B.; Chen, Y.; and Neville, J., eds., *Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI 2023, Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence, IAAI 2023, Thirteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2023, Washington, DC, USA, February 7-14, 2023*, 14304–14311. AAAI Press.
- Klipfel, A.; Fregier, Y.; Sayede, A.; and Bouraoui, Z. 2023b. Unified Model For Crystalline Material Generation. In *Proceedings of the International Joint Conference on Artificial Intelligence, IJCAI*.
- Long, T.; Fortunato, N. M.; Opahle, I.; Zhang, Y.; Samathrakakis, I.; Shen, C.; Gutfleisch, O.; and Zhang, H. 2021. Constrained crystals deep convolutional generative adversarial network for the inverse design of crystal structures. *npj Computational Materials*, 7(1): 66.
- Luo, Y.; Yan, K.; and Ji, S. 2021. GraphDF: A Discrete Flow Model for Molecular Graph Generation. In Meila, M.; and Zhang, T., eds., *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, 7192–7203. PMLR.
- Noh, J.; Kim, J.; Stein, H. S.; Sanchez-Lengeling, B.; Gregoire, J. M.; Aspuru-Guzik, A.; and Jung, Y. 2019. Inverse

- Design of Solid-State Materials via a Continuous Representation. *Matter*, 1(5): 1370–1384.
- Nouira, A.; Sokolovska, N.; and Crivello, J.-C. 2019. CrystalGAN: Learning to Discover Crystallographic Structures with Generative Adversarial Networks. arXiv:1810.11203.
- Pathak, Y.; Juneja, K. S.; Varma, G.; Ehara, M.; and Priyakumar, U. D. 2020. Deep learning enabled inorganic material generator. *Phys. Chem. Chem. Phys.*, 22: 26935–26943.
- Ren, Z.; Tian, S. I. P.; Noh, J.; Oviedo, F.; Xing, G.; Li, J.; Liang, Q.; Zhu, R.; Aberle, A. G.; Sun, S.; Wang, X.; Liu, Y.; Li, Q.; Jayavelu, S.; Hippalgaonkar, K.; Jung, Y.; and Buonassisi, T. 2022. An invertible crystallographic representation for general inverse design of inorganic crystals with targeted properties. *Matter*, 5(1): 314–335.
- Rusman, N.; and Dahari, M. 2016. A review on the current progress of metal hydrides material for solid-state hydrogen storage applications. *International Journal of Hydrogen Energy*, 41(28): 12108–12126.
- Satorras, V. G.; Hoogeboom, E.; and Welling, M. 2021. E(n) Equivariant Graph Neural Networks. arXiv:2102.09844.
- Sawada, Y.; Morikawa, K.; and Fujii, M. 2019. Study of Deep Generative Models for Inorganic Chemical Compositions. arXiv:1910.11499.
- Schneuing, A.; Du, Y.; Harris, C.; Jamasb, A.; Igashov, I.; Du, W.; Blundell, T.; Lió, P.; Gomes, C.; Welling, M.; Bronstein, M.; and Correia, B. 2022. Structure-based Drug Design with Equivariant Diffusion Models. arXiv:2210.13695.
- Shi\*, C.; Xu\*, M.; Zhu, Z.; Zhang, W.; Zhang, M.; and Tang, J. 2020. GraphAF: a Flow-based Autoregressive Model for Molecular Graph Generation. In *International Conference on Learning Representations*.
- Sohl-Dickstein, J.; Weiss, E.; Maheswaranathan, N.; and Ganguli, S. 2015. Deep Unsupervised Learning using Nonequilibrium Thermodynamics. In Bach, F.; and Blei, D., eds., *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, 2256–2265. Lille, France: PMLR.
- Tournet, J.; Lee, Y.; Karuturi, S. K.; Tan, H. H.; and Jagdish, C. 2020. III–V Semiconductor Materials for Solar Hydrogen Production: Status and Prospects. *ACS Energy Letters*, 5(2): 611–622.
- Trippe, B. L.; Yim, J.; Tischer, D.; Baker, D.; Broderick, T.; Barzilay, R.; and Jaakkola, T. S. 2023. Diffusion Probabilistic Modeling of Protein Backbones in 3D for the motif-scaffolding problem. In *The Eleventh International Conference on Learning Representations*.
- Wu, J.; Shen, T.; Lan, H.; Bian, Y.; and Huang, J. 2021. SE(3)-Equivariant Energy-based Models for End-to-End Protein Folding. *bioRxiv*.
- Xie, T.; Fu, X.; Ganea, O.-E.; Barzilay, R.; and Jaakkola, T. 2021. Crystal Diffusion Variational Autoencoder for Periodic Material Generation. arXiv:2110.06197.
- Xie, T.; and Grossman, J. C. 2018. Crystal Graph Convolutional Neural Networks for an Accurate and Interpretable Prediction of Material Properties. *Physical Review Letters*, 120(14).
- Xu, M.; Yu, L.; Song, Y.; Shi, C.; Ermon, S.; and Tang, J. 2022. GeoDiff: A Geometric Diffusion Model for Molecular Conformation Generation. In *International Conference on Learning Representations*.
- Yan, K.; Liu, Y.; Lin, Y.; and Ji, S. 2022. Periodic Graph Transformers for Crystal Material Property Prediction. In *The 36th Annual Conference on Neural Information Processing Systems*.
- Zhao, Y.; Al-Fahdi, M.; Hu, M.; Siriwardane, E. M. D.; Song, Y.; Nasiri, A.; and Hu, J. 2021. High-Throughput Discovery of Novel Cubic Crystal Materials Using Deep Generative Neural Networks. *Advanced Science*, 8(20): 2100566.
- Zhao, Y.; Siriwardane, E. M. D.; Wu, Z.; Fu, N.; Al-Fahdi, M.; Hu, M.; and Hu, J. 2023. Physics guided deep learning for generative design of crystal materials with symmetry constraints. *npj Computational Materials*, 9(1): 38.