# PLAID: PREFERENCE ALIGNED LANGUAGE MODEL FOR TARGETED INORGANIC MATERIALS DESIGN

**Andy Xu, Rohan Desai, Larry Wang, Gabriel Hope, Ethan Ritz**
Harvey Mudd College
Claremont, CA 91711, USA
{andxu, rodesai, larwang, ghope, eritz}@hmc.edu

## ABSTRACT

Discovering novel materials is critical for technological advancements such as solar cells, batteries, and carbon capture. However, the development of new materials is constrained by a slow and expensive trial-and-error process. To accelerate this pipeline, we introduce PLaID, a Large Language Model (LLM) fine-tuned for stable crystal generation. We first fine-tune a base version of LLaMA-2 7B on Wyckoff-based text representations of crystals. Then, we further fine-tune via Direct Preference Optimization on sampled structures categorized by their stability. By encoding symmetry constraints directly into text and aligning model outputs to explore stable chemical space, PLaID generates structures that are thermodynamically stable, unique, and novel at a 40% higher rate than prior methods. Our work demonstrates the potential of adapting post-training techniques from natural language processing to materials design, paving the way for targeted and efficient discovery of novel materials.

## 1 INTRODUCTION

The discovery of new solid-state materials is the foundation of many transformative technologies including solar cells (Green et al., 2014), batteries (Zhao et al., 2020), and carbon capture (Sriram et al., 2024a). However, the search for new materials is constrained by the immense scale of chemical space—previous explorations have only uncovered a fraction of the total number of potential stable inorganic compounds (Davies et al., 2016). Generative models offer a promising avenue for accelerating technological breakthroughs by efficiently discovering novel and stable structures in unexplored regions of chemical space.

Previous work have applied variational autoenconders (Xie et al., 2021), denoising models (Zeni et al., 2023; Jiao et al., 2023; Miller et al., 2024), and language models (Gruver et al., 2024; Flam-Shepherd & Aspuru-Guzik, 2023) to generate stable and novel structures. Language models are advantageous due to their natural language prompting, which allows for a single model to be applied to many tasks like conditional generation, infilling, and crystal structure prediction. Specifically, the pre-trained knowledge of large language models (LLMs) make them data-efficient, requiring fewer training examples compared to training models from scratch (Gruver et al., 2024).

Symmetry is a defining aspect of crystal structures. The set of rotations, reflections, inversions, and translations exhibited by a crystal lattice form its *space group*. A crystal's space group and symmetries are not merely a mathematical construct but critical to many optical, electrical, and magnetic properties like piezoelectricity (Malgrange et al., 2014; Yang et al., 2005). One way to define a crystal's structure from its symmetries is via Wyckoff positions, whereby one can specify a few key atomic coordinates and have the remaining atomic positions filled by applying symmetry operations (Hahn et al., 1983).

Recent methods like WyFormer (Kazeev et al., 2024) and CrystFormer (Cao et al., 2024) have incorporated symmetry into language models for crystal generation. However, both works use bespoke chemical tokenization and require explicit *a priori* space group and formula constraints, limiting their utility in *de novo* generation. Other fields of machine learning have shown that general-purpose architectures that focus on scale and data expressivity outperform models with handcrafted, domain-specific constraints (Dosovitskiy, 2020; Kaplan et al., 2020). We contend that the same "Bitter

Lesson" holds true for materials generation (Sutton, 2019). Rather than imposing explicit crystallographic constraints, we train an LLM to learn and exploit structural parameters from the data itself. By allowing the model to discover patterns implicitly through training on Wyckoff-based text representations, we see significant improvements in unconditional generation.

To guide our search space towards chemically useful structures, we apply Direct Preference Optimization (DPO) (Rafailov et al., 2023), a method from Reinforcement Learning from Human Feedback (RLHF) used to align LLMs with human preferences (Bai et al., 2022). Our pipeline is as follows: first, we perform supervised finetuning on a base LLM with Wyckoff-based text encodings of crystals, then further fine-tune the LLM via DPO on generated structures categorized by their stability. By incorporating inherent crystal symmetries and the feedback of a neural network interatomic potential, we significantly increase the rate of stable, unique, and novel (S.U.N.) materials. Our experiments demonstrate that our method **PLaID generates stable materials at $\sim$60% higher rate and S.U.N materials at over a 40% higher rate** than prior models while retaining the flexibility of natural language prompting.

Our contributions are as follows:

1. We develop a novel, symmetry-informed text representation for crystal structures which is compact, performant, and physically-motivated.
2. We apply DPO to materials generation and show its ability to guide behavior towards desirable chemical properties.
3. We demonstrate that PLaID significantly outperforms existing state-of-the-art models in generating novel and stable materials.

## 2 METHOD

### 2.1 BASE MODEL: FINE-TUNING LLAMA FOR CRYSTAL GENERATION

Our approach begins with a pre-trained large language model (LLaMA-2) (Touvron et al., 2023), which has demonstrated strong capabilities in text generation and reasoning. Following the methodology outlined by (Gruver et al., 2024), we fine-tune the base version of LLaMA-2 7B to model the crystal generation process via parameter-efficient fine-tuning (Hu et al., 2021). Unlike prior works that rely on explicit space group constraints or predefined templates, our method enables more flexible and generalizable unconditional structure generation through the use of a natural language prompt.

The training objective is to maximize the likelihood that the model generates the correct token sequence representing stable crystal structures. We fine-tune two versions of the base model: one where the crystal structures are represented as text-based 3D coordinates and another based on a new text-based Wyckoff representation.

### 2.2 WYCKOFF REPRESENTATION: ENCODING SYMMETRY INTO TEXT

Crystal structures inherently follow symmetry constraints that can be represented by space groups and Wyckoff positions. To improve the model's ability to generate valid and stable structures, we introduce a **Wyckoff representation** that encodes these symmetries directly into text. Unlike previous methods such as DiffCSP++ (Jiao et al., 2024) which rely on searching template structures for Wyckoff positions, our approach allows direct prediction of these sites within the generative process.

Because Wyckoff positions specify atomic site symmetries in a given space group, they significantly increase the possibility of feasible atomic arrangements. Our encoding method represents each crystal as a sequence of tokens as shown in Figure 1. We specify the chemical formula, space group, lattice parameters, individual chemical elements, fractional coordinates, and Wyckoff site symmetries. Note that we include the chemical formula first to ensure self-consistency during generation.

### 2.3 DIRECT PREFERENCE OPTIMIZATION (DPO) FOR STABILITY ALIGNMENT

While encoding symmetry improves structural validity, stability remains a crucial factor for practical material discovery. To explicitly align our model toward generating more stable crystals, we employ

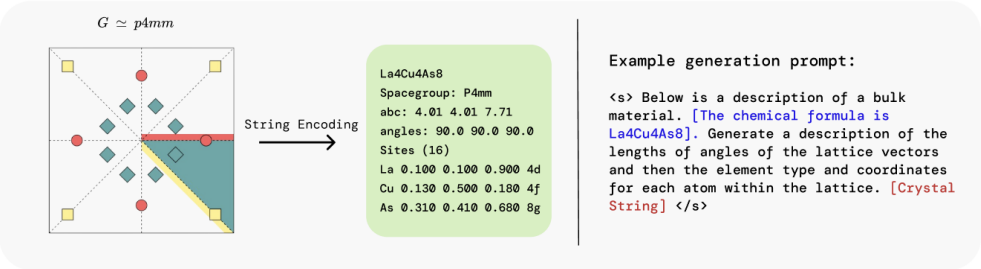


Figure 1: **Left:** An example crystal highlighting the $p4mm$ space group symmetry, where colors represent atoms of different elements. By leveraging symmetry, the asymmetric unit of the crystal (highlighted in green) can be used to represent the entire crystal in our Wyckoff-based text representation. **Right:** An example prompt used during training. The blue conditioning information is optional, and the crystal string is replaced with the encoding on the left.

**Direct Preference Optimization (DPO)** (Rafailov et al., 2023), a reinforcement learning technique that directly optimizes model outputs based on ranked preferences.

To train the model to prefer stable crystals, we first construct a crystal stability preferences dataset using the **M3GNet relaxation model**. M3GNet is a machine learning interatomic potential (MLIP) that predicts relaxed formation energies for materials, serving as an efficient proxy for stability assessment. To generate this dataset, we sample 10,000 crystal structures from our fine-tuned LLaMA model. From here, we categorize crystals based on metastability as predicted by M3GNet ($< 0.08$ eV/atom $E^{\text{hull}}$). Specifically, we create stability preference pairs $(x, y_w, y_l)$, where $x$ is the model prompt and $y_w$ and $y_l$ are accepted and rejected crystals.

We apply DPO to fine-tune LLaMA on this curated preference dataset. DPO is advantageous compared to other reinforcement learning methods because it eliminates the need for a computationally-expensive reward model, directly leveraging pairwise comparisons for model alignment. Given a pre-trained reference model $\pi_{\text{ref}}$ which assigns a probability to our prompt $x$, the loss can be written as

$$\mathcal{L}_{\text{DPO}} = -\mathbb{E}_{(x,y_w,y_l)}\left[\log\sigma\left(\beta\log\frac{\pi_\theta(y_w\,|\,x)}{\pi_{\text{ref}}(y_w\,|\,x)} - \beta\log\frac{\pi_\theta(y_l\,|\,x)}{\pi_{\text{ref}}(y_l\,|\,x)}\right)\right] \tag{1}$$

where $\pi_\theta$ is the model to be fine-tuned, and $\beta$ is a hyper-parameter controlling how far the tuned model deviates from the reference model. We have additional information about our DPO hyperparameters and dataset creation process in Appendix A.2.

## 3 Experiments

### 3.1 Setup

We trained our model on the well-established MP-20 dataset (Xie et al., 2021), a collection of 45,231 inorganic crystalline materials from the Materials Project (Jain et al., 2013). The dataset includes structures with up to 20 atoms, all of which are metastable. We follow the methodology of Gruver et al. (2024) by independently fine-tuning a pre-trained LLaMA-2 7B model using 4-bit quantization and Low-Rank Adapters (LoRA) (Hu et al., 2021), implemented with PyTorch (Paszke et al., 2019) and Transformers (Wolf et al., 2020). Symmetry information for our Wyckoff representation is calculated by Pyxtal Fredericks et al. (2021). Following supervised fine-tuning, we apply DPO on the generated preference dataset to further guide generation towards stable structures. Full hyperparameters details are provided in Appendix A.2. We sample 10,000 structures from each fine-tuned model, parsing a CIF from the generated string. We resample if a CIF cannot be parsed from the string, which guarantees all samples can be interpreted as crystals but does not guarantee validity. Additional information relating to sampling validity and efficiency are available in Appendix A.2.

Table 1: Results for unconditional materials generation on the MP-20 dataset. Shows ablated results for PLaID, where base refers to LLaMA-2 7B using a coordinate representation and no DPO.

| Method | Params | Validity (%) (↑) | | Coverage (%) (↑) | | Property Dist. (↓) | | Stability (%) (↑) | | S.U.N. (%) (↑) |
|---|---|---|---|---|---|---|---|---|---|---|
| | | Structural | Composition | Recall | Precision | $d_\rho$ | $d_{\text{elem}}$ | Initial | Relaxed | |
| CDVAE | – | 99.93 | 86.93 | 98.31 | 99.35 | 0.914 | 1.654 | 0.1 | 3.6 | 3.5 |
| DiffCSP | – | 99.61 | 82.23 | 99.53 | 99.35 | 0.257 | 0.403 | 8.9 | 12.5 | 9.7 |
| DiffCSP++ | – | **99.99** | 85.81 | 99.48 | **99.66** | 0.278 | 0.408 | 8.9 | 13.2 | 9.1 |
| FlowMM | – | 96.43 | 83.37 | 99.47 | 99.71 | 0.291 | **0.079** | 4.1 | 9.3 | 6.3 |
| SymmCD | – | 94.32 | 85.85 | **99.64** | 98.87 | **0.090** | 0.399 | 5.0 | 9.4 | 7.0 |
| CrystalLLM 7B† | – | 96.4 | **93.3** | 91.1 | 94.9 | 3.61 | 1.06 | – | – | – |
| **LLaMA-2 7B** | | | | | | | | | | |
| PLaID-Types | Base | 99.21 | 87.98 | 97.56 | 99.55 | 0.824 | 0.084 | 6.3 | 17.0 | 9.2 |
| | DPO | 98.40 | 88.91 | 96.99 | 99.56 | 1.325 | 0.165 | 7.9 | **20.9** | 12.2 |
| | Wyckoff | 99.87 | 91.28 | 98.32 | 99.64 | 0.129 | 0.110 | **11.4** | 17.5 | 12.0 |
| PLaID | Wyckoff & DPO | 99.67 | 93.03 | 98.07 | 99.44 | 1.500 | 0.102 | 8.6 | 20.5 | **14.2** |

† Our CrystalLLM reproduction, PLaID Base, significantly outperforms reported CrystalLLM results. We explore this in Appendix A.2.

### 3.2 METRICS

To initially evaluate the quality of our generated crystal structures, we focus on **validity**, **coverage**, and **property statistics** as defined by Xie et al. (2021). These metrics provide an effective proxy for assessing the quality and diversity of generated crystals before conducting more computationally expensive stability evaluations. We explain more details about these metrics in Appendix A.1

Our primary metric for evaluating generated crystal structures is the **S.U.N Rate** from Miller et al. (2024), which measures the percentage of crystals that are **stable, unique**, and **novel**. Stability is assessed by comparing a crystal's energy to a convex hull of previously computed energies from Riebesell et al. (2023). Crystals on or below the hull ($\leq 0$ eV/atom $E^{\text{hull}}$) are deemed *stable*. Due to compute constraints, we assess stability via the CHGNet MLIP (Deng et al., 2023). Similar to SymmCD (Levy et al., 2024), we randomly sub-sample 1,000 crystals of the 10,000 samples and predict their stability before and after CHGNet relaxation. Note that we specifically **use different MLIPs**—M3GNet and CHGNet—for preference dataset creation and crystal generation evaluation to avoid reward hacking or overfitting to a specific MLIP. Following Miller et al. (2024), we evaluate **uniqueness**—differentiation from other generated crystals—and **novelty**, which measures diversity from the training data to calculate S.U.N.

### 3.3 RESULTS

We compare our model to five prior methods, as reported from Levy et al. (2024): CDVAE (Xie et al., 2021), DiffCSP (Jiao et al., 2023), FlowMM (Miller et al., 2024), SymmCD (Levy et al., 2024) and CrystalLLM (Gruver et al., 2024). Our main results are presented in Table 1. On stability and S.U.N, the most important metrics, PLaID significantly outperforms all prior methods. While both Wyckoff text encoding and DPO fine-tuning each individually improved the base model's stability and S.U.N rates, applying both methods together significantly improved the rates, resulting in our best PLaID model. We note that the pre-relaxation stability for both DPO fine-tuned models goes down, however we theorize this may be due to the quality of relaxations from the M3GNet pseudo-potential model. For our best overall PLaID model, 20.5% of the relaxed structures are deemed stable, out of which 74% are unique and 94% are novel, leading to a S.U.N. rate of 14.20%. **PLaID achieves a $\sim$60% higher stability rate and 40% higher S.U.N rate than the best previous method.**

On the proxy metrics, PLaID achieves performance on par with other methods, and reaches the highest compositional validity. Our Wyckoff encoding specifically increases compositional validity for both the base and DPO fine-tuned models, highlighting symmetry-based encodings as a natural and intuitive mechanism to increase the validity and stability of generated structures. Though many of these metrics have saturated, we report PLaID's performance for comparison and completeness.

## 4 DISCUSSION

In this paper, we introduce a novel Wyckoff-based text encoding for crystal structures and demonstrate DPO's ability to guide generation towards chemically stable and useful structures. Our method achieves state-of-the-art performance across stability and S.U.N metrics. Current random structure search methods achieve less than a 1% success rate (Pickard & Needs, 2011) in identifying stable materials—PLaID represents a significant acceleration from traditional approaches to future methods with greater stability and real-world utility.

In future work, we plan to (1) apply DPO to conditional generation tasks like generating crystals with desired electronic and magnetic properties, (2) explore how scaling laws dictate generation performance, and (3) verify our results via Density Functional Theory Kohn & Sham (1965).

### ACKNOWLEDGMENTS

This work used high performance computing resources from Accelerating Computing For Emerging Sciences (ACES) at Texas A&M University through allocation CIS240657 from the Advanced Cyberinfrastructure Coordination Ecosystem: Services Support (ACCESS) program, which is supported by U.S. National Science Foundation grants 2138259, 2138286, 2138307, 2137603, and 2138296.

## REFERENCES

Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*, 2022.

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

Daniel W Davies, Keith T Butler, Adam J Jackson, Andrew Morris, Jarvist M Frost, Jonathan M Skelton, and Aron Walsh. Computational screening of all stoichiometric inorganic materials. *Chem*, 1(4):617–627, 2016.

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

Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.

Robert A Evarestov and Vyacheslav P Smirnov. *Site symmetry in crystals: theory and applications*. Springer Berlin Heidelberg, 1997.

Daniel Flam-Shepherd and Alán Aspuru-Guzik. Language models can generate molecules, materials, and protein binding sites directly in three dimensions as xyz, cif, and pdb files. *arXiv preprint arXiv:2305.05708*, 2023.

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

Martin A Green, Anita Ho-Baillie, and Henry J Snaith. The emergence of perovskite solar cells. *Nature photonics*, 8(7):506–514, 2014.

Nate Gruver, Anuroop Sriram, Andrea Madotto, Andrew Gordon Wilson, C. Lawrence Zitnick, and Zachary Ward Ulissi. Fine-tuned language models generate stable inorganic materials as text. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=vN9fpfqoP1.

Theo Hahn, Uri Shmueli, and JC Wilson Arthur. *International tables for crystallography*, volume 1. Reidel Dordrecht, 1983.

Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models, 2021.

Anubhav Jain, Shyue Ping Ong, Geoffroy Hautier, Wei Chen, William Davidson Richards, Stephen Dacek, Shreyas Cholia, Dan Gunter, David Skinner, Gerbrand Ceder, and Kristin A. Persson. Commentary: The materials project: A materials genome approach to accelerating materials innovation. *APL Materials*, 1(1), 7 2013. doi: 10.1063/1.4812323.

Rui Jiao, Wenbing Huang, Peijia Lin, Jiaqi Han, Pin Chen, Yutong Lu, and Yang Liu. Crystal structure prediction by joint equivariant diffusion. *Advances in Neural Information Processing Systems*, 36:17464–17497, 2023.

Rui Jiao, Wenbing Huang, Yu Liu, Deli Zhao, and Yang Liu. Space group constrained crystal generation. *arXiv preprint arXiv:2402.03992*, 2024.

Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.

Nikita Kazeev, Ruiming Zhu, Ignat Romanov, Andrey E Ustyuzhanin, Shuya Yamazaki, Wei Nong, and Kedar Hippalgaonkar. Wyckofftransformer: Generation of symmetric crystals. In *AI for Accelerated Materials Design-NeurIPS 2024*, 2024.

Walter Kohn and Lu Jeu Sham. Self-consistent equations including exchange and correlation effects. *Physical review*, 140(4A):A1133, 1965.

Daniel Levy, Siba Smarak Panigrahi, Sékou-Oumar Kaba, Qiang Zhu, Mikhail Galkin, Santiago Miret, and Siamak Ravanbakhsh. Symmcd: Symmetry-preserving crystal generation with diffusion models. In *AI for Accelerated Materials Design-NeurIPS 2024*, 2024.

Cécile Malgrange, Christian Ricolleau, and Michel Schlenker. *Symmetry and physical properties of crystals*. Springer, 2014.

Benjamin Kurt Miller, Ricky TQ Chen, Anuroop Sriram, and Brandon M Wood. Flowmm: Generating materials with riemannian flow matching. *arXiv preprint arXiv:2406.04713*, 2024.

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

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library, 2019. URL `http://arxiv.org/abs/1912.01703`. cite arxiv:1912.01703Comment: 12 pages, 3 figures, NeurIPS 2019.

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

Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36, 2023.

Janosh Riebesell, Rhys EA Goodall, Anubhav Jain, Philipp Benner, Kristin A Persson, and Alpha A Lee. Matbench discovery–an evaluation framework for machine learning crystal stability prediction. *arXiv preprint arXiv:2308.14920*, 2023.

Anuroop Sriram, Sihoon Choi, Xiaohan Yu, Logan M Brabson, Abhishek Das, Zachary Ulissi, Matt Uyttendaele, Andrew J Medford, and David S Sholl. The open dac 2023 dataset and challenges for sorbent discovery in direct air capture. *ACS Central Science*, 10(5):923–941, 2024a. doi: 10.1021/acscentsci.3c01629. URL https://doi.org/10.1021/acscentsci.3c01629.

Anuroop Sriram, Benjamin Kurt Miller, Ricky T. Q. Chen, and Brandon M Wood. FlowLLM: Flow matching for material generation with large language models as base distributions. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024b. URL https://openreview.net/forum?id=0bFXbEMz8e.

Richard Sutton. The bitter lesson. *Incomplete Ideas (blog)*, 13(1):38, 2019.

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.

Logan Ward, Ankit Agrawal, Alok Choudhary, and Christopher Wolverton. A general-purpose machine learning framework for predicting properties of inorganic materials. *npj Computational Materials*, 2(1):1–7, 2016.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. Transformers: State-of-the-art natural language processing. In Qun Liu and David Schlangen (eds.), *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pp. 38–45, Online, October 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-demos.6. URL https://aclanthology.org/2020.emnlp-demos.6/.

Tian Xie, Xiang Fu, Octavian-Eugen Ganea, Regina Barzilay, and Tommi Jaakkola. Crystal diffusion variational autoencoder for periodic material generation. *arXiv preprint arXiv:2110.06197*, 2021.

Jiashi Yang et al. *An introduction to the theory of piezoelectricity*, volume 9. Springer, 2005.

Claudio Zeni, Robert Pinsler, Daniel Zügner, Andrew Fowler, Matthew Horton, Xiang Fu, Sasha Shysheya, Jonathan Crabbé, Lixin Sun, Jake Smith, et al. Mattergen: a generative model for inorganic materials design. *arXiv preprint arXiv:2312.03687*, 2023.

Claudio Zeni, Robert Pinsler, Daniel Zügner, Andrew Fowler, Matthew Horton, Xiang Fu, Zilong Wang, Aliaksandra Shysheya, Jonathan Crabbé, Shoko Ueda, et al. A generative model for inorganic materials design. *Nature*, pp. 1–3, 2025.

Qing Zhao, Sanjuna Stalin, Chen-Zi Zhao, and Lynden A Archer. Designing solid-state electrolytes for safe, energy-dense batteries. *Nature Reviews Materials*, 5(3):229–252, 2020.

Nils ER Zimmermann and Anubhav Jain. Local structure order parameters and site fingerprints for quantification of coordination environment and crystal structure similarity. *RSC advances*, 10 (10):6063–6081, 2020.

# A Appendix

## A.1 Proxy Metric Details

**Validity** provides a computationally efficient check on whether a generated crystal is physically plausible. We assess this through two criteria: *structural validity*, which ensures that no two atoms are closer than 0.5 Å, and *positional validity*, which verifies charge neutrality.

**Coverage** measures how well the generated structures represent the diversity of real crystals. We compute recall and precision by comparing the CrystalNN (Zimmermann & Jain, 2020) and Magpie (Ward et al., 2016) fingerprints of generated samples against those in the test set, using their pairwise distances.

**Property statistics** evaluate the alignment of generated crystals with the distribution of real materials. We compare the generated and test set crystals by computing the Wasserstein distances of key properties, specifically atomic density ($d_\rho$) and the number of unique elements ($d_{elem}$), to quantify distributional similarity.

**Stability Rate** assesses how many generated materials are thermodynamically stable. Stability is determined by the energy above the hull ($E^{\text{hull}}$) metric, which measures the energy difference between a material and the convex hull of competing phases with the same composition. A material is considered stable if $E^{\text{hull}} \leq 0$ eV/atom, while those with $E^{\text{hull}} < 0.08$ eV/atom are classified as metastable. We compute $\text{E}_{\text{hull}}$ by relaxing generated structures using CHGNet using the Materials Project dataset (February 2023) as a reference. Total energies were corrected using the MP2020 compatability scheme, which maintains consistency across various functionals (DFT/DFT+U).

**S.U.N. Rate** extends the stability metric to assess both novelty and uniqueness. A structure is novel if it is not structurally similar to any training set material, determined using Pymatgen's StructureMatcher (Ong et al., 2013). Uniqueness ensures that duplicate generations are not counted separately by grouping structurally similar outputs into equivalence classes. The S.U.N. rate is defined as the fraction of generated structures that are Stable, Unique, and Novel:

$$\text{Stability Rate} = \frac{N_{\text{stable}}}{N_{\text{gen}}} \times 100\% \tag{2}$$

$$\text{SUN Rate} = \frac{N_{\text{SUN}}}{N_{\text{gen}}} \times 100\% \tag{3}$$

Together, these metrics provide a more rigorous assessment of generative model performance by quantifying both the stability and originality of discovered materials.

### A.2 Additional Experiment Details

**Training/Evaluation Pipeline.** Our finetuning pipeline consists of multiple stages, including the supervised finetuning of a base LLaMA-2 7B model, direct preference optimization (DPO), and evaluation using CHGNet's relaxation model. The complete process is illustrated in Figure 2.

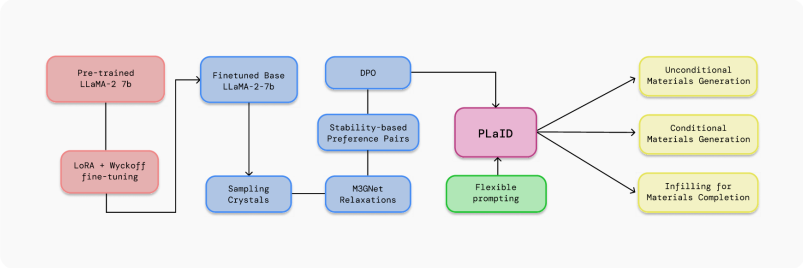


Figure 2: Flowchart of the model finetuning pipeline. Note that although unconditional generation was the focus of this work, we support infilling and conditional materials generation.

**Supervised Finetuning.** We perform supervised finetuning (SFT) on LLaMA to adapt the model for crystal structure generation. For our fine-tuning, the model performs infilling one-third of the time. The remaining two-thirds of the time, we generate structures *de novo*, where we either generate crystals unconditionally, or randomly append between one and five additional properties in our prompt. The properties we include are the chemical formula, energy above the hull, formation energy per atom, band gap and space group.

To fine-tune different LLaMA-2 7B models on both text-based and Wyckoff-based crystal representation, we followed the steps outlined by Gruver et al. (2024) to reproduce their results. We use an AdamW optimizer with batch size of 16 samples, a learning rate of $10^{-5}$, fp-4 mixed precision alongside LoRA adapters (with a LoRA rank of 8, LoRA alpha of 32, and LoRA dropout of 0.05) to fine-tune over the MP-20 dataset for 10 epochs.

While running the SFT experiments, we noticed an interesting phenomenon. Sticking to Gruver et al. (2024)'s procedure caused the model to encounter gradient-NaN errors in the training process. By switching the precision of the model weights not included in LoRA to bfloat-16 (as opposed to float-16), our gradients not only became computable but the base fine-tuned model performance also vastly exceeded that of the fine-tuned LLaMA-2 7B outlined in Bai et al. (2022). We hypothesize that this may be because the LlaMA-2 models were trained using bfloat-16, making it more numerically stable to maintain the same precision format during fine-tuning. This alignment in numerical precision likely helps preserve the model's learned representations and computational stability, leading to more effective parameter updates during the fine-tuning process.

**DPO.** After supervised finetuning, we apply Direct Preference Optimization (DPO) to align the model's probability distribution to favor more stable structures. We use the DPO Trainer from TRL (Wolf et al., 2020) using the Adam optimizer with a batch size of 16 samples, a learning rate of $10^{-6}$, fp-4/bfloat-16 precision, and a $\beta$ of $0.1$ on one epoch of our dataset.

We curate the dataset used for DPO by sampling $10,000$ structures from the fine-tuned LLaMA-2 7B models. From these structures, we filter the crystals into two categories falling above or below the $0.08$ eV/atom $E^{\text{hull}}$ threshold. For DPO, those considered metastable constitute our accepting crystals and those considered non-metastable constitute our rejecting crystals in our preference pairs.

We form our dataset by pairing accepted and rejected crystals in a 1:3 ratio determined experimentally via internal ablation studies (that is every meta-stable crystal gets paired with 3 non meta-stable crystals). This creates roughly 10,000 to 20,000 DPO pairs depending on the metastability rate. Interestingly, the model incurs significantly worse performance when we compile our dataset using a 1:1 and 1:10 ratio. We hypothesize that smaller ratios perform worse as our training dataset becomes noticeably smaller, and larger ratios perform worse as the model overfits to each positive sample and consequently generates less diverse structures and lower S.U.N rates.

**CHGNet + M3GNet.** We employ M3GNet to curate our preference dataset (assessing crystal structures above and below the 0.08 eV/atom $E^{\text{hull}}$ threshold) and CHGNet for our stability evaluations by computing relaxed formation energies. For M3GNet and CHGNet, we used 500 and 1500 relaxation steps respectively.

**LLM Sampling Efficiency.** We benchmark PLaID sampling efficiency on a singular NVIDIA H100. We generate the maximum batch size we can fit on a singular H100 GPU of 256, and find PLaID takes approximately 88.8 minutes to generate 10,000 samples. We find that roughly 75% of generated strings are parsable as crystals. A more useful metric is the time to generate a S.U.N. material, which takes the sampling time divided by the number of S.U.N. materials categorized by CHGNet. Here, PLaID takes 3.75 seconds seconds per S.U.N. material, which is competitive with other generative materials models (Miller et al., 2024; Sriram et al., 2024b).

Though it is not the focus of this work, we notice that the time to generate samples decreases to roughly 20 minutes and the parsing rate of generated strings increases to over 95% as we scale to the larger Alex-MP-20 dataset (Zeni et al., 2025).

**Experiments – Server Details.** Training was conducted on a high performance computing (HPC) cluster equipped with NVIDIA H100 GPUs. We used a single NVIDIA H100 GPU for LLaMA-2 7B SFT and DPO as well as inference. We used an AMD Epyc 7443P for S.U.N. rate evaluations.

### A.3 Additional Results

We compare visualizations between the space group distribution of the MP-20 dataset we used to perform SFT on our base models, a sample of 10,000 crystals we generate from our PLaID model trained on DPO finetuning and Wyckoff representations, and a sample of 10,000 crystals we generate from the baseline model that has only undergone SFT on the 3D coordinate representation in Figure 3.

From the diagram, it appears clear that the DPO finetuning process does not alter the model's generated space group distribution in a noticeable way. We observe that DPO fine-tuning is significantly overestimating the frequency of the most common groups (space groups *Fm-3m* (225), *Pnma* (62), and *P63/mmc* (194)). We suspect this behavior is a result of bias from the original LLM, and that

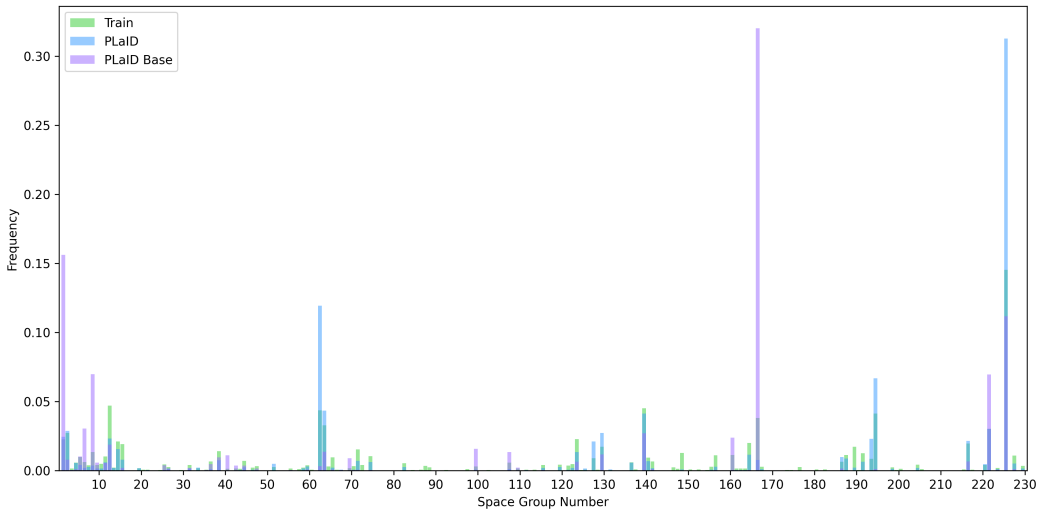


Figure 3: A comparison between the three space group distributions allows us to extract valuable insights on how our PLaID phases adjust model behavior.

the fine-tuning is exploiting regions with more data. In the future, a possible avenue of exploration is to introduce a regularizer to combat this.

Meanwhile, space groups *Cm* (8) and *P1* (1) are prominent in the base PLaID model crystals. Space group *Cm* has a single symmetry operation (a glide plane), while space group *P1* only has the identity operation. Notably, both are among the five most frequently generated space groups in the model, even though they are not in the top five of the training set. This supports the hypothesis that the Wyckoff representation helps PLaID generate systems with complex symmetry requirements, skewing the distribution towards generating crystals with more symmetries and ensuring a more even distribution across all space groups.

### A.4 Miscellaneous Examples of Generated Crystals

We present additional qualitative examples of crystal structures generated by our model. These structures are visualized using the ASE and PyMatGen toolkits. Figure 4 show several samples from different space groups.

### A.5 Wyckoff Theory: Mathematical Background

A Wyckoff position defines a set of equivalent atomic sites in a given space group (Evarestov & Smirnov, 1997). These positions are characterized by symmetry constraints, which reduce the degrees of freedom in atomic placement.

For a crystal structure belonging to a space group $G$, the Wyckoff positions can be defined as:

$$W = \{gx | g \in G\} \tag{4}$$

where $x$ is a fractional atomic coordinate, and $g$ is an element of the space group.

Each Wyckoff site follows symmetry constraints that ensure atoms occupy positions dictated by the underlying crystallographic symmetry:

$$x' = Rx + t \tag{5}$$

where $R$ is a symmetry operation (such as a rotation, reflection, or inversion) that can be represented as a matrix, and $t$ is a translation vector.

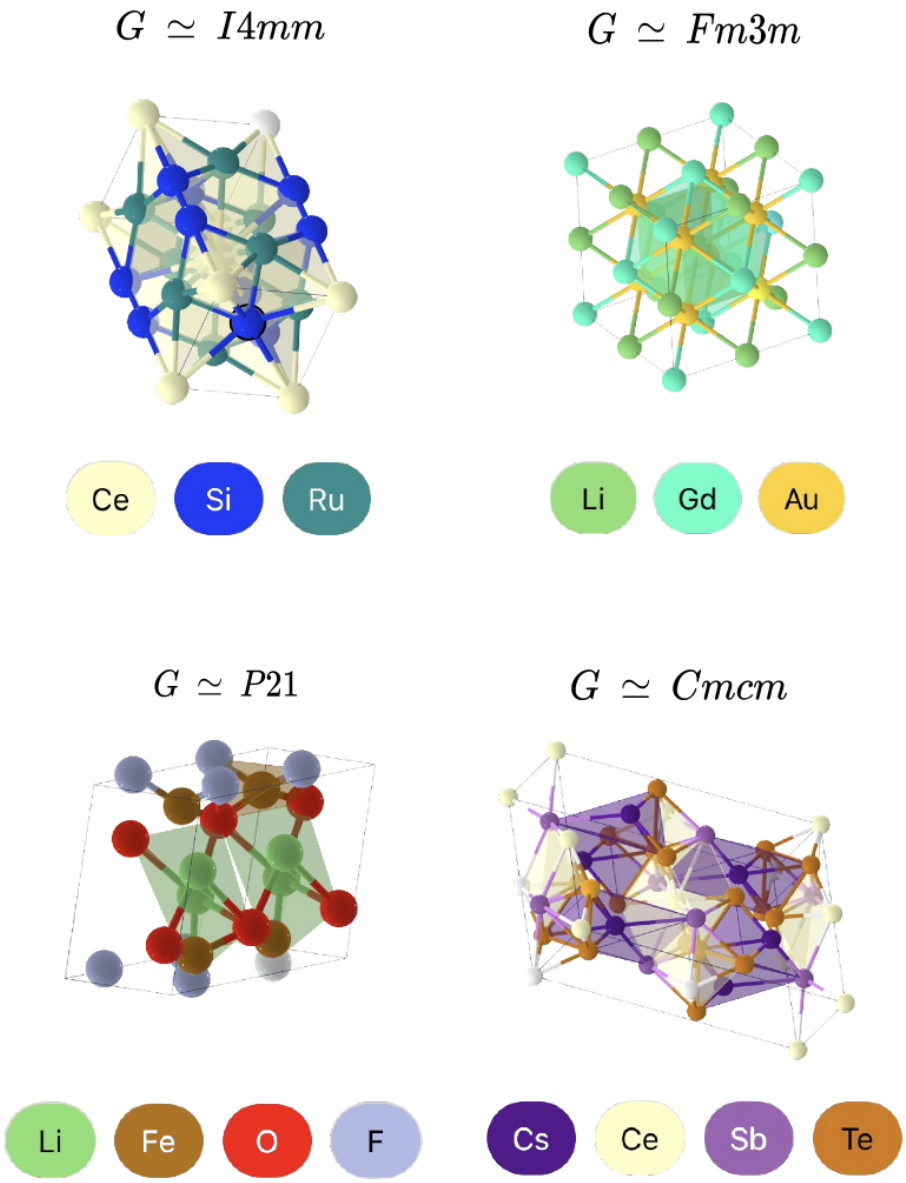


Figure 4: Examples of crystals generated by our LLaMA-2-7B model fine-tuned with Wyckoff representation and Direct Preference Optimization (DPO). The crystals' corresponding chemical formulas are from top left to bottom right respectively, $Ce_2Si_4Ru_4$ , $Li_4Gd_4Au_8$, $Cs_4Ce_4Sb_4Te_{12}$, and $Li_4Fe_4O_4F_4$, respectively.

In this notation, $x'$ is considered to be part of the same crystallographic orbit as $x$, where $x$ is referred to as the generating point. As any one arbitrary point in the crystallographic orbit can be used to generate all other points in the orbit, the choice of $x$ is not unique.

All symmetry operations $(R, t)$ which map a point in the crystallographic orbit to itself form a finite subgroup of $G$, which can be identified as one of the 32 crystallographic point groups. Each Wyckoff site in each space group is associated with one such point group, and is often labeled according to the multiplicity of $W$ (the number of points in the orbit generated by the space group in a single unit cell) and a letter encoding the point group (i.e. '4a').