
The Loss Landscape of XRD-Based Structure Optimization Is Too Rough for Gradient Descent

Nofit Segal¹, Akshay Subramanian¹, Mingda Li², Benjamin Kurt Miller³, and Rafael Gómez-Bombarelli¹

¹Materials Science and Engineering, MIT

²Nuclear Science and Engineering, MIT

³FAIR at Meta

{nofit, akshay_s, mingda, rafagb}@mit.edu, bkmi@meta.com

Abstract

Solving crystal structures from powder X-ray diffraction (XRD) is a central challenge in materials characterization. One machine learning approach to inverting XRD involves optimizing a loss function that compares a ground-truth XRD spectrum against an XRD from a generated candidate crystal. We investigate the roughness of this optimization landscape by comparing reference structures with physically motivated distortions. We show that XRD similarity metrics result in a highly non-convex landscape, complicating direct optimization by gradient descent. Constraining the optimization symmetrically within the ground-truth crystal family improves recovery; nevertheless, the landscape may remain non-convex along symmetry axes. While energy-based relaxation excels at locating low-energy configurations, it cannot directly target the specific minima implied by spectra. We therefore advocate incorporating symmetry inductive biases directly into XRD-conditioned generative models, followed by energy relaxation, to enable more reliable reconstruction of phases from diffraction patterns.

1 Introduction

Determining the atomic structure of a crystal from its powder X-ray diffraction (XRD) pattern is a longstanding and central challenge in materials characterization [1, 2]. The inverse problem of recovering the full three-dimensional crystal structure solely from an XRD pattern is fundamentally inaccessible due to the loss of phase information in the scattered waves [3][4, Chapter 9.3, 13.2]. Nevertheless, powder diffraction is widely used for identifying and characterizing crystalline solids. In practice, this is typically achieved by comparing the observed XRD spectrum to a reference database and performing least-squares refinement, known as Rietveld analysis [5–7]. However, this process is highly sensitive to initial parameters [8], and more importantly, it relies on the presence of the correct structure in the database and cannot be used to reconstruct novel or unreported phases.

Experimental phenomena such as preferred orientation, peak overlap, crystal twinning, and instrumental noise further complicate structural determination from powder XRD patterns [9–11]. Moreover, many minerals and metallic alloys exhibit solid solution ranges with very slight lattice shifts, which result in ranges of stoichiometries with nearly the same diffraction pattern [4, Chapter 10.3.2][12–14], making XRD-to-structure mapping a one-to-many problem. Consequently, accurate structure reconstruction typically requires refinement model fitting and domain-specific prior knowledge.

From a computational perspective, structural ambiguity remains even under idealized conditions. Two structures with different compositions can exhibit highly correlated XRD patterns if they share similar symmetry [15–17]. Moreover, even when stoichiometry is fixed, structures with close though

distinct space groups can yield highly similar XRD patterns [2]. Notably, small distortions in lattice parameters or atomic coordinates can cause discontinuous changes in the diffraction pattern, such as the appearance or disappearance of peaks due to shifting Bragg conditions [4, Chapter 9]. This introduces a highly non-smooth relationship between structure and XRD signal.

Recently, there has been a surge of interest in crystal structure determination from XRD patterns using generative modeling [18–23]. A growing body of work applies gradient-based optimization approaches that leverage differentiable physics to refine generated or otherwise-obtained crystal structures by minimizing the difference between simulated and target XRD patterns. For example, Riesel et al. [19] use a differentiable XRD simulator to update both lattice parameters and atomic positions via gradient descent (GD). Parackal et al. [24] assume the lattice parameters are known and restrict the optimization to atomic positions along Wyckoff degrees of freedom. Bayesian optimization has also been used to morph candidate structures by maximizing the cosine similarity between the XRD patterns [25]. GD has further been applied to determine lattice parameters through indexing single-crystal diffraction patterns [26].

Complementary efforts have focused on developing metrics for evaluating the similarity of diffraction patterns. Otero-de-la Roza [27] introduced a cross-correlation-based metric that captures equivalence between diffraction patterns while remaining invariant to lattice distortions. Building on this work, Racioppi et al. [28] applied the metric to crystal structure prediction from XRD data, jointly optimizing the structure by minimizing both this similarity metric and the structure’s enthalpy. Hernández-Rivera et al. [29] systematically analyzed the sensitivity of different families of similarity metrics under isotropic lattice strain. Li et al. [30] proposed an entropy-based similarity measure for spectra and demonstrated its utility for molecular database retrieval from mass spectrometry data.

In this work, we explore the powder XRD-to-structure mapping through the lens of GD optimization. The goal is to recover correct structures based solely on XRD similarity. While the ground truth structure is available for evaluation, in practice, it is unknown, reflecting the challenge of reconstructing novel materials solely from diffraction data. We introduce two types of physically motivated distortions, random lattice distortions and uncorrelated atomic displacements, to emulate experimentally relevant effects such as thermal expansion from lattice vibrations and thermal fluctuations [8, 31–33]. These distorted structures also resemble outputs from generative models for crystal structure prediction, which often produce nearly correct geometries but with imperfect symmetry [34–36].

We find that mapping XRD patterns to crystal structures is challenging because high diffraction agreement, as currently measured in literature, does not ensure structural accuracy. We show that commonly used XRD similarity measures, such as cosine similarity, Mean Squared Error (MSE), and Entropy similarity, are sensitive to both lattice and coordinate noise distortions, and optimization between distorted structures and ground-truth XRD diffraction can become trapped in local minima. We explore an alternative strategy that enforces symmetry constraints, highlighting the role of symmetry in connecting XRD to the crystal structure.

2 Method

We selected 10 structures from the MP20 dataset [38], a collection of small, inorganic, thermodynamically (meta)stable structures from the Materials Project [39]. We selected according to the most common space groups (see appendix 4), which span a range of crystal symmetries: $P6/mmm$, $Pn\bar{3}m$, $I\bar{4}$, Cm , $I\bar{4}/mmm$, $Fm\bar{3}m$, $C2/m$, $P6_3/mmc$, $Pm\bar{3}m$, Pm . For each structure, we generated 50 distorted versions using two noise models:

Lattice noise. Random lattice distortions were applied via strain tensors [11]. These modifications alter the cell while keeping fractional atomic coordinates fixed. Each distorted structure was generated by randomly sampling the entries of a strain tensor and applying it to the ground-truth lattice matrix. Let $\mathbf{L} \in \mathbb{R}^{3 \times 3}$ be the ground-truth lattice matrix (columns are the lattice vectors). We noise the lattice by a random deformation matrix $\mathbf{S} \in \mathbb{R}^{3 \times 3}$,

$$\tilde{\mathbf{L}} = \mathbf{S} \mathbf{L}, \quad \mathbf{S} = \begin{pmatrix} s_{11} & s_{12} & s_{13} \\ s_{21} & s_{22} & s_{23} \\ s_{31} & s_{32} & s_{33} \end{pmatrix},$$

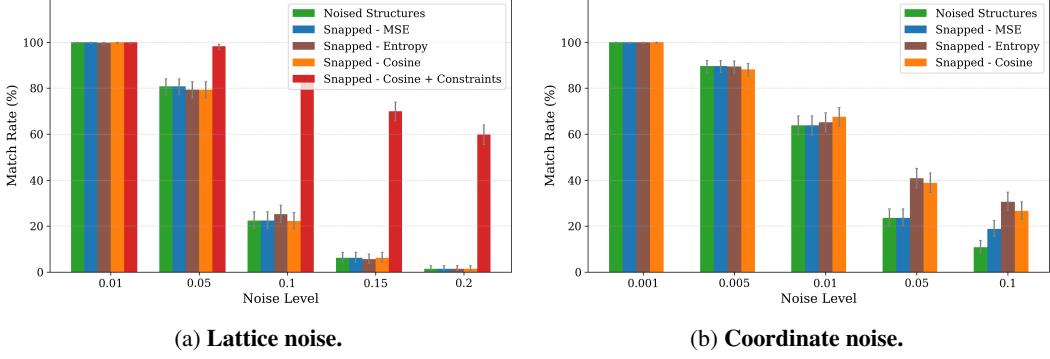


Figure 1: **Results of XRD-based optimization under two types of structural noise.** Crystal structures were optimized with respect to XRD similarity metrics using the `snaps` method [19], which struggles to recover the correct structure under both lattice and coordinate perturbations. The plots show match rates computed with `StructureMatcher` ($l_{tol} = 0.1$, $stol = 0.2$, $angle_tol = 5^\circ$) under random lattice (a) and coordinate (b) perturbations. Error bars represent 95% Jeffreys binomial credible intervals [37]. For lattice distortions, incorporating symmetry constraints significantly improves robustness, even at high noise levels.

while keeping atomic fractional coordinates fixed. For a noise level $\sigma_l > 0$, the entries of \mathbf{S} are sampled as

$$\begin{aligned} s_{ii} &\sim \text{Unif}(1 - \sigma_l, 1 + \sigma_l), & i \in \{1, 2, 3\}, \\ s_{ij} &\sim \text{Unif}(-\sigma_l, \sigma_l), & i \neq j. \end{aligned}$$

Thus, diagonal entries produce uniaxial expansion/compression, whereas off-diagonal entries induce shear. By construction, these distortions do not preserve crystal symmetry and can introduce diverse deformation modes. If \mathbf{f} denotes a fractional coordinate, then the Cartesian position changes from $\mathbf{x} = \mathbf{L}\mathbf{f}$ to $\tilde{\mathbf{x}} = \tilde{\mathbf{L}}\mathbf{f}$ with \mathbf{f} unchanged.

Coordinate noise. Independent, uncorrelated positional perturbations were applied to each atom by adding Gaussian-distributed noise to its fractional coordinates. Let $x_i^{(n)} \in [0, 1]^3$ be the fractional coordinates of atom n . For a noise scale $\sigma_c > 0$, we draw i.i.d. perturbations

$$\varepsilon_n \sim \mathcal{N}(0, \sigma_c^2 I_3), \quad n = 1, \dots, N,$$

and set the noisy fractional coordinates to

$$x_f^{(n)} = w(x_i^{(n)} + \varepsilon_n), \quad w(u) = u - \lfloor u \rfloor \in [0, 1]^3,$$

where $\lfloor u \rfloor$ applies the floor function componentwise. Equivalently, $x_f^{(n)} \equiv x_i^{(n)} + \varepsilon_n \pmod{1}$ (elementwise), i.e., on the 3-torus $\mathbb{T}^3 = \mathbb{R}^3 / \mathbb{Z}^3$.

After applying noise, we attempted to recover the ground-truth structure via gradient-based optimization using `StructSnap`, a differentiable XRD simulator [19]. We compute diffraction patterns from the structure factor contributions of each atomic site, following Bragg’s law and the kinematic scattering model. The resulting pattern is a 2D tensor of 2θ angles and intensities.

Optimization. We refine either the lattice parameters or atomic coordinates in accordance with the noise applied. The structure is passed through this differentiable diffraction pipeline to produce a simulated pattern, which is then compared to the ground-truth pattern using a chosen loss function. Given distorted-state XRD $\hat{\mathbf{x}}$ and ground-truth XRD \mathbf{x} (see A.2), we minimize the negative cosine similarity, MSE loss, or negative entropy similarity, which are defined respectively:

$$\begin{aligned} \mathcal{L}_{\text{cos}} &= \frac{1}{N} \sum_{i=1}^N \left(-\frac{\mathbf{x}^{(i)} \cdot \hat{\mathbf{x}}^{(i)}}{\|\mathbf{x}^{(i)}\|_2 \|\hat{\mathbf{x}}^{(i)}\|_2} \right), & \mathcal{L}_{\text{MSE}} &= \frac{1}{N} \sum_{i=1}^N \left\| \mathbf{x}^{(i)} - \hat{\mathbf{x}}^{(i)} \right\|_2^2, \\ \mathcal{L}_{\text{entropy}} &= -\frac{1}{N} \sum_{i=1}^N \left(1 - \frac{2S_{\hat{\mathbf{x}}\mathbf{x}} - S_{\mathbf{x}} - S_{\hat{\mathbf{x}}}}{\log 4} \right). \end{aligned}$$

where S denotes the Shannon entropy. Gradients of the loss are backpropagated to update structural parameters using PyTorch’s autograd.

Symmetry Constraints. We examine the effect of symmetry-based constraints in the lattice-noise case by enforcing the ground-truth crystal family during optimization. Although this information may not be available in practice, recent work has shown it can be classified from XRD data [40–42]. Using projected optimization, each gradient step is followed by projection onto the constrained values.

Let $\theta = (a, b, c, \alpha, \beta, \gamma)$ be the lattice parameters. Simple constrained gradient descent on a proposed crystal $\hat{\mathcal{X}}$ using user-selected loss $\mathcal{L} \in \{\mathcal{L}_{\text{cos}}, \mathcal{L}_{\text{MSE}}, \mathcal{L}_{\text{entropy}}\}$ and symmetry projection operator \mathcal{P} would be:

$$\theta^0 = \mathcal{P}(\theta^{\text{init}}), \quad \theta^{k+1} = \mathcal{P} \left(\theta^k - \nabla_{\theta^k} \mathcal{L} \left(\text{xrd} \left(\hat{\mathcal{X}}(\theta^k) \right), \text{xrd}_0 \right) \right),$$

with the iterations repeating until some convergence criterion is satisfied. Note that θ^{init} are initial lattice parameters from our prediction, model, or, in this case, distorted ground-truth; $\hat{\mathcal{X}}(\theta^k)$ is the representation of our proposed crystal structure, which depends on current lattice parameters θ^k ; xrd is a map from crystal to computed powder x-ray diffraction pattern; xrd_0 is the xrd pattern of the reference we aim to recover. The projection operator \mathcal{P} is defined by relevant crystal family, e.g.,

$$\mathcal{P}_{\text{cubic}}(a, b, c, \alpha, \beta, \gamma) = (\bar{a}, \bar{a}, \bar{a}, 90^\circ, 90^\circ, 90^\circ), \quad \bar{a} = \frac{a+b+c}{3}.$$

Thus, a , b , and c are first updated independently according to their gradients, then set to the mean value \bar{a} , while angles are fixed to 90° . Similar projectors are defined for the remaining crystal families (see A.3).

Recovery performance is assessed primarily using Match Rate, the fraction of optimized structures that are identified as structurally equivalent to the ground-truth by `StructureMatcher` [43], considering lattice, atomic positions, and symmetry. The tolerances used are 0.1 for lattice, 0.2 for atomic site positions, and 5 degrees for angles.

3 Results

We performed optimization on distorted crystal structures across a range of noise types and levels. For each condition, 50 distorted versions were generated for each of 10 ground truth structures, yielding 500 distorted inputs per noise setting. The 50 variations per structure enable a statistical view. We report the match rate of the optimized structures to the ground truth in Figure 1.

Figure 1 illustrates that XRD-based optimization degrades significantly as noise increases. In particular, for lattice distortions, the largest drop occurs between noise levels of 0.05 and 0.1. While 0.1 is the lattice tolerance threshold for matching (see Section 2), the noise level defines the maximum possible strain sampled, so lattice lengths remain within the tolerable range. Using either cosine similarity, MSE loss or entropy similarity as the objective makes only small differences in performance.

3.1 Symmetry Constraints: Strengths and Limitations

Incorporating symmetry-based constraints during XRD-based optimization notably improves robustness to lattice noise for many structures in this study, as shown in Figure 1a by the higher match rates achieved when constraints are applied. Our constraints (see subsection 2) project updates back into the correct crystal family at each optimization step, thereby guiding the search along a reduced-dimensionality symmetry-consistent path and helping the optimizer avoid local minima unrelated to the desired symmetry.

Figures 2, 3 and 5 illustrate the non-convex nature of the XRD-based loss landscape with respect to lattice parameters, through **simplified 2D cross-sections** of the optimization landscape. In Figure 2, we distort the lattice parameters a and c of U_2Ti , a hexagonal structure of space group $P6/mmm$ (No. 191), and compute the cosine similarity loss between distorted structures’ spectra and the ground truth. The resulting contour map reveals multiple deep local minima, indicating the optimizer’s potential to get trapped in suboptimal solutions. The three most prominent local minima are highlighted, and their corresponding XRD patterns are shown in the right panel. Despite their structural deviation

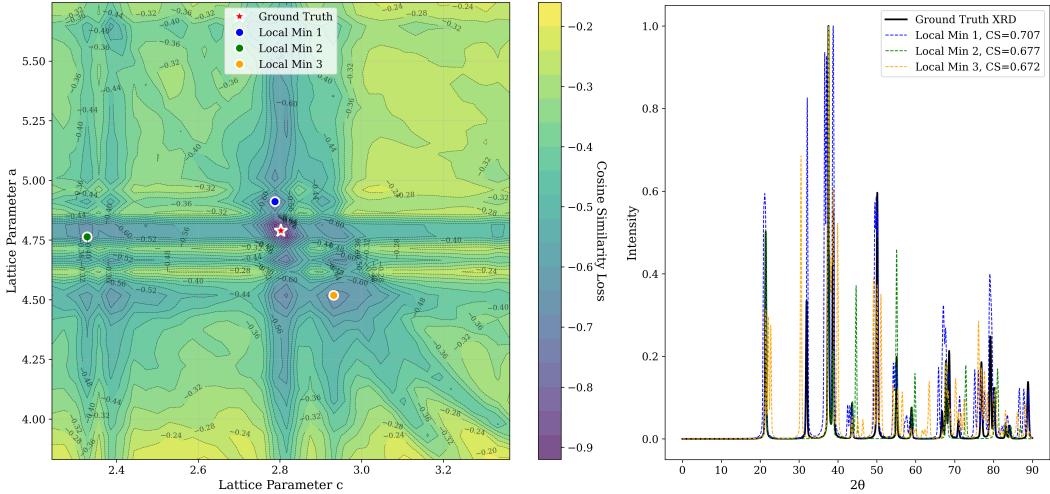


Figure 2: **2D landscape of XRD cosine similarity (CS) loss as a function of lattice parameters a and c of U_2Ti structure, illustrating the presence of multiple local minima.** (left) Cosine similarity loss topographic map showing non-convex behavior with several local minima. (right) XRD patterns for the structures corresponding to the marked local minima: all exhibit reasonably high cosine similarity to the ground truth pattern despite having different lattice parameters.

from the true lattice parameters, the patterns show high cosine similarity to the ground truth due to subtle shifts and peak splittings that preserve the overall spectral profile.

Figure 3 illustrates how symmetry constraints facilitate accurate structure reconstruction. For clarity, only a two-dimensional slice of the optimization landscape is shown, although the full optimization occurs in a six-dimensional space for distorted lattices. Simulated GD trajectories are visualized for two representative structures under three settings: (i) unconstrained GD, (ii) unconstrained GD initialized at a constrained point, and (iii) fully constrained GD.

In Figure 3a, the lattice parameters a and α of Au_2S (cubic, space group $Pn\bar{3}m$, No. 224) are perturbed. The unconstrained GD converges to a distant local minimum, while constrained initialization improves convergence but remains suboptimal. The fully constrained trajectory successfully recovers the ground truth. In Figure 3b, we distort a and γ of $Na_3MnCoNiO_6$, (monoclinic, space group Cm , No. 8). Here, unconstrained GD, even from a constrained initialization ($\gamma = 90^\circ$), fails to reach the ground truth, whereas fully constrained GD converges correctly. Although initialization points were selected for illustration, similar behavior holds across broader regions. These examples highlight how symmetry-constrained XRD optimization improves convergence and yields higher match rates than unconstrained methods (Figure 1).

Notably, in Figure 5, we observe fluctuations that pose challenges for symmetry-constrained GD along symmetry axes such as $a = b$ and $\alpha = 90^\circ$. While fluctuations along $a = b$ are pronounced, those along $\alpha = 90^\circ$ are comparatively shallow and may be mitigated through techniques such as momentum [44] or regularization, which were not studied in this work. Although symmetry constraints generally improve refinement performance, the landscape visualized here highlights that XRD-based GD may remain sensitive to initialization and prone to local minima in some cases.

Figure 6 illustrates a case where high XRD pattern cosine similarity does not imply structural similarity. After applying noise, the structure no longer matches the ground truth, with the diffraction pattern exhibiting substantial peak shifts and new reflections. The unconstrained GD optimizer converges to a local minimum where many peaks align with those of the ground truth but correspond to different Miller indices hkl , meaning they arise from different atomic planes, indicating a distinct underlying structure. Despite this, the XRD pattern achieves a deceptively high cosine similarity of 0.71. When crystal-family constraints are imposed during optimization, the optimizer returns a structure that matches the ground truth within the defined tolerance. The resulting diffraction pattern shows slightly shifted peaks, reflecting differences in lattice parameters. As a result, the cosine similarity drops to just 0.05. The convergence to this shallow minimum is likely driven by fluctuations

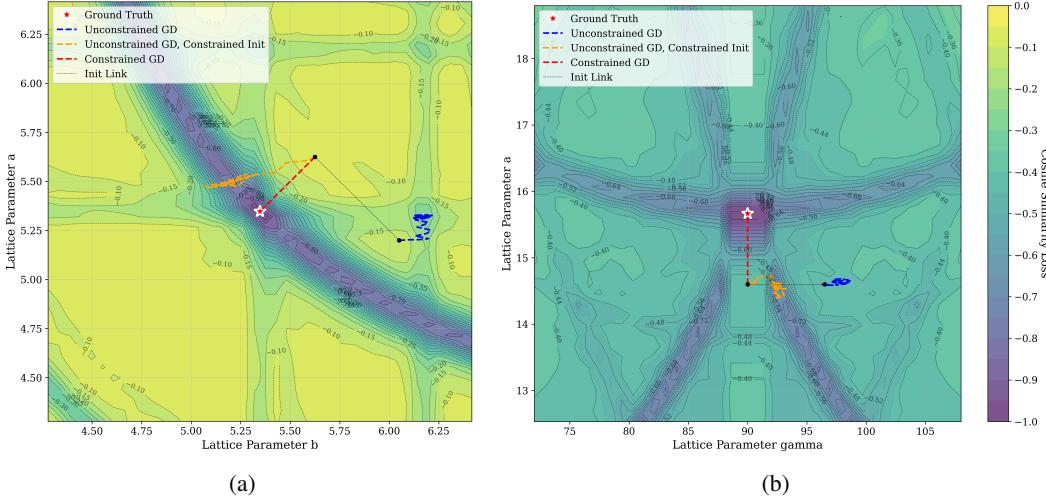


Figure 3: **2D landscape of XRD cosine similarity (CS) loss as a function of lattice parameters, with simulated optimization paths for XRD-based gradient descent (GD).** Unconstrained GD, unconstrained GD with a constrained initialization, and fully constrained GD. Unconstrained GD converges to some local minima, even with constrained initialization, whereas constrained GD reaches the ground truth. (a) Lattice parameters a and b of cubic Au_2S are perturbed. (b) Lattice parameters a and γ of monoclinic $\text{Na}_3\text{MnCoNiO}_6$ are perturbed.

along the symmetry axis. This counterintuitive outcome highlights a key limitation of using XRD pattern similarity as the sole reconstruction objective: similar structures can yield dissimilar patterns, and conversely, distinct structures may appear similar, under such metrics.

3.2 Comparison to Energy Relaxation

Structural relaxation through potential energy minimization is often used [45–49] to refine candidate structures. Figure 7a shows that the universal ML interatomic potential (MLIP) CHGNet [45] accurately recovers structures matching the ground truth from the distorted state alone, except for a high level (0.1) of coordinate noise. This suggests that the loss landscape of energy relaxation is much smoother than XRD similarity [50]. Cross-sections of the energy optimization landscape are shown in Figures 7b, showcasing a smooth and convex behavior.

While MLIPs accurately reconstruct structures under small and controlled distortions as in this study, their optimization landscape is not globally smooth. In the context of inverse XRD, the goal is to recover a **specific structure** from its diffraction pattern, corresponding to a **particular minimum** on the potential energy surface. Large distortions can displace the system into the basin of attraction of a different local minimum, as evident for coordinate perturbations of 10% (Figure 7b).

4 Conclusion

XRD provides a direct experimental link for generative crystal modeling, enabling the identification of novel phases. Our results highlight symmetry’s role in bridging XRD and structure, but also reveal that in some cases the XRD-to-structure landscape may remain non-convex even along symmetry axes, making post-hoc optimization difficult. We illustrate this with physically motivated random distortions, though generative models may introduce more complex biases. The observations in this work suggest that progress in inverse XRD will require new generative architectures that condition on XRD and embed symmetry as an inductive bias, with final refinements guided by energy relaxation.

5 Acknowledgments

N.S.'s research is sponsored by the Google-Schwarzman College of Computing Fellowship and the National Science Foundation (NSF) under Award Number 2209892 'Garden: A FAIR Framework for Publishing and Applying AI Models for Translational Research in Science, Engineering, Education, and Industry'.

References

- [1] William Lawrence Bragg. The analysis of crystals by the x-ray spectrometer. *Proceedings of the Royal society of London. Series A, Containing papers of a mathematical and physical character*, 89(613):468–489, 1914.
- [2] William IF David and Kenneth Shankland. Structure determination from powder diffraction data. *Foundations of Crystallography*, 64(1):52–64, 2008.
- [3] Herbert A Hauptman. The phase problem of x-ray crystallography. *Reports on Progress in Physics*, 54(11):1427, 1991.
- [4] Christopher Hammond. *The basics of crystallography and diffraction*, volume 21. Oxford university press, 2015.
- [5] Hugo M Rietveld. A profile refinement method for nuclear and magnetic structures. *Journal of applied Crystallography*, 2(2):65–71, 1969.
- [6] Stacy Gates-Rector and Thomas Blanton. The powder diffraction file: a quality materials characterization database. *Powder diffraction*, 34(4):352–360, 2019.
- [7] FH Allen and VJ Hoy. The cambridge structural database (csd). In *International Tables for Crystallography Volume F: Crystallography of biological macromolecules*, pages 663–668. Springer, 2006.
- [8] CM Biwer, Z Feng, D Finstad, M McDonnell, M Knezevic, M McKerns, DJ Savage, and SC Vogel. Spotlight: efficient automated global optimization in rietveld analysis of diffraction data. *Scientific Reports*, 15(1):8358, 2025.
- [9] Cameron F Holder and Raymond E Schaak. Tutorial on powder x-ray diffraction for characterizing nanoscale materials, 2019.
- [10] N Chandra, K Ravi Acharya, and PCE Moody. Analysis and characterization of data from twinned crystals. *Biological Crystallography*, 55(10):1750–1758, 1999.
- [11] Nathan J Szymanski, Sean Fu, Ellen Persson, and Gerbrand Ceder. Integrated analysis of x-ray diffraction patterns and pair distribution functions for machine-learned phase identification. *Npj computational materials*, 10(1):45, 2024.
- [12] Huaming Deng, Kaijun Wang, Yunbiao Duan, Weijun Zhang, and Jin Hu. Preparation of in/sn nanoparticles (in₃sn and insn₄) by wet chemical one-step reduction and performance study. *Coatings*, 12(4):429, 2022.
- [13] Qin Wang, Shanshan Cai, Shiyu Yang, Yongjian Yu, Yongkang Wan, Jubo Peng, Jiajun Wang, and Xiaojing Wang. Comparison of high-speed shear properties of low-temperature sn-bi/cu and sn-in/cu solder joints. *Journal of Materials Science: Materials in Electronics*, 35(8):576, 2024.
- [14] AM El-Taher, SA Mansour, and IH Lotfy. Robust effects of in, fe, and co additions on microstructures, thermal, and mechanical properties of hypoeutectic sn–zn-based lead-free solder alloy. *Journal of Materials Science: Materials in Electronics*, 34(7):599, 2023.
- [15] Shuai Gu, Bitian Fu, Gjergj Dodbiba, Toyohisa Fujita, and Baizeng Fang. A sustainable approach to separate and recover indium and tin from spent indium–tin oxide targets. *RSC Advances*, 7(82):52017–52023, 2017.

[16] A Ayeshamariam, M Kashif, M Bououdina, U Hashim, M Jayachandran, and ME Ali. Morphological, structural, and gas-sensing characterization of tin-doped indium oxide nanoparticles. *Ceramics International*, 40(1):1321–1328, 2014.

[17] M Ashokkumar, K Palanisamy, A Ganesh Kumar, C Muthusamy, and KJ Senthil Kumar. Green synthesis of silver and copper nanoparticles and their composites using ocimum sanctum leaf extract displayed enhanced antibacterial, antioxidant and anticancer potentials. *Artificial Cells, Nanomedicine, and Biotechnology*, 52(1):438–448, 2024.

[18] Gabe Guo, Tristan Luca Saidi, Maxwell W Terban, Michele Valsecchi, Simon JL Billinge, and Hod Lipson. Ab initio structure solutions from nanocrystalline powder diffraction data via diffusion models. *Nature Materials*, pages 1–9, 2025.

[19] Eric A Riesel, Tsach Mackey, Hamed Nilforoshan, Minkai Xu, Catherine K Badding, Alison B Altman, Jure Leskovec, and Danna E Freedman. Crystal structure determination from powder diffraction patterns with generative machine learning. *Journal of the American Chemical Society*, 146(44):30340–30348, 2024.

[20] Frederik Lizak Johansen, Ulrik Friis-Jensen, Erik Bjørnager Dam, Kirsten Marie Ørnsbjerg Jensen, Rocío Mercado, and Raghavendra Selvan. decifer: Crystal structure prediction from powder diffraction data using autoregressive language models. *arXiv preprint arXiv:2502.02189*, 2025.

[21] Qi Li, Rui Jiao, Liming Wu, Tiannian Zhu, Wenbing Huang, Shifeng Jin, Yang Liu, Hongming Weng, and Xiaolong Chen. Powder diffraction crystal structure determination using generative models. *Nature Communications*, 16(1):7428, 2025.

[22] Qingsi Lai, Fanjie Xu, Lin Yao, Zhifeng Gao, Siyuan Liu, Hongshuai Wang, Shuqi Lu, Di He, Liwei Wang, Linfeng Zhang, et al. End-to-end crystal structure prediction from powder x-ray diffraction. *Advanced Science*, 12(8):2410722, 2025.

[23] Gabe Guo, Judah Goldfeder, Ling Lan, Aniv Ray, Albert Hanming Yang, Boyuan Chen, Simon JL Billinge, and Hod Lipson. Towards end-to-end structure determination from x-ray diffraction data using deep learning. *npj Computational Materials*, 10(1):209, 2024.

[24] Abhijith S Parackal, Rhys EA Goodall, Felix A Faber, and Rickard Armiento. Identifying crystal structures beyond known prototypes from x-ray powder diffraction spectra. *Physical Review Materials*, 8(10):103801, 2024.

[25] Joohwi Lee, Junpei Oba, Nobuko Ohba, and Seiji Kajita. Creation of crystal structure reproducing x-ray diffraction pattern without using database. *npj Computational Materials*, 9(1):135, 2023.

[26] Yaroslav Gevorkov, Oleksandr Yefanov, Anton Barty, Thomas A White, Valerio Mariani, Wolfgang Brehm, Aleksandra Tolstikova, R-R Grigat, and Henry N Chapman. Xgandalf—extended gradient descent algorithm for lattice finding. *Foundations of Crystallography*, 75(5):694–704, 2019.

[27] Alberto Otero-de-la Roza. Powder-diffraction-based structural comparison for crystal structure prediction without prior indexing. *Applied Crystallography*, 57(5), 2024.

[28] Stefano Racioppi, Alberto Otero-de-la Roza, Samad Hajinazar, and Eva Zurek. Powder x-ray diffraction assisted evolutionary algorithm for crystal structure prediction. *Digital Discovery*, 4(1):73–83, 2025.

[29] Efraín Hernández-Rivera, Shawn P Coleman, and Mark A Tschopp. Using similarity metrics to quantify differences in high-throughput data sets: application to x-ray diffraction patterns. *ACS combinatorial science*, 19(1):25–36, 2017.

[30] Yuanyue Li, Tobias Kind, Jacob Folz, Arpana Vaniya, Sajjan Singh Mehta, and Oliver Fiehn. Spectral entropy outperforms ms/ms dot product similarity for small-molecule compound identification. *Nature Methods*, 18(12):1524–1531, 2021.

[31] Oliviero Cannelli, Julia Wiktor, Nicola Colonna, Ludmila Leroy, Michele Puppin, Camila Bacellar, Ilia Sadykov, Franziska Krieg, Grigory Smolentsev, Maksym V Kovalenko, et al. Atomic-level description of thermal fluctuations in inorganic lead halide perovskites. *The Journal of Physical Chemistry Letters*, 13(15):3382–3391, 2022.

[32] Federico Brivio, Jarvist M Frost, Jonathan M Skelton, Adam J Jackson, Oliver J Weber, Mark T Weller, Alejandro R Goni, Aurélien MA Leguy, Piers RF Barnes, and Aron Walsh. Lattice dynamics and vibrational spectra of the orthorhombic, tetragonal, and cubic phases of methylammonium lead iodide. *Physical Review B*, 92(14):144308, 2015.

[33] L Delgado-Aparicio, M Bitter, Y Podpaly, J Rice, W Burke, M Sanchez Del Rio, P Beiersdorfer, R Bell, R Feder, C Gao, et al. Effects of thermal expansion of the crystal lattice on x-ray crystal spectrometers used for fusion research. *Plasma Physics and Controlled Fusion*, 55(12):125011, 2013.

[34] Rui Jiao, Wenbing Huang, Yu Liu, Deli Zhao, and Yang Liu. Space group constrained crystal generation. In *The Twelfth International Conference on Learning Representations*, 2024.

[35] Daniel Levy, Siba Smarak Panigrahi, Sékou-Oumar Kaba, Qiang Zhu, Kin Long Kelvin Lee, Mikhail Galkin, Santiago Miret, and Siamak Ravanbakhsh. Symmcld: Symmetry-preserving crystal generation with diffusion models. *arXiv preprint arXiv:2502.03638*, 2025.

[36] Nikita Kazeev, Wei Nong, Ignat Romanov, Ruiming Zhu, Andrey Ustyuzhanin, Shuya Yamazaki, and Kedar Hippalgaonkar. Wyckoff transformer: Generation of symmetric crystals. *arXiv preprint arXiv:2503.02407*, 2025.

[37] Lawrence D Brown, T Tony Cai, and Anirban DasGupta. Interval estimation for a binomial proportion. *Statistical science*, 16(2):101–133, 2001.

[38] Tian Xie, Xiang Fu, Octavian-Eugen Ganea, Regina Barzilay, and Tommi S Jaakkola. Crystal diffusion variational autoencoder for periodic material generation. In *International Conference on Learning Representations*, 2021.

[39] Anubhav Jain, Shyue Ping Ong, Geoffroy Hautier, Wei Chen, William Davidson Richards, Stephen Dacek, Shreyas Cholia, Dan Gunter, David Skinner, Gerbrand Ceder, et al. Commentary: The materials project: A materials genome approach to accelerating materials innovation. *APL materials*, 1(1), 2013.

[40] Yuta Suzuki, Hideitsu Hino, Takafumi Hawai, Kotaro Saito, Masato Kotsugi, and Kanta Ono. Symmetry prediction and knowledge discovery from x-ray diffraction patterns using an interpretable machine learning approach. *Scientific reports*, 10(1):21790, 2020.

[41] Bin Cao, Yang Liu, Zinan Zheng, Rui Feng Tan, Jia Li, and Tong-yi Zhang. Simxrd-4m: Big simulated x-ray diffraction data accelerate the crystal symmetry classification. *arXiv preprint arXiv:2406.15469*, 2024.

[42] Byung Do Lee, Jin-Woong Lee, Woon Bae Park, Joonseo Park, Min-Young Cho, Satendra Pal Singh, Myoung Ho Pyo, and Kee-Sun Sohn. Powder x-ray diffraction pattern is all you need for machine-learning-based symmetry identification and property prediction. *Advanced Intelligent Systems*, 4(7):2200042, 2022.

[43] 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.

[44] Ning Qian. On the momentum term in gradient descent learning algorithms. *Neural networks*, 12(1):145–151, 1999.

[45] 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.

- [46] Ilyes Batatia, Philipp Benner, Yuan Chiang, Alin M Elena, David P Kovacs, Janosh Riebesell, Xavier R Advincula, Mark Asta, Matthew Avaylon, William J Baldwin, et al. A foundation model for atomistic materials chemistry. *arXiv preprint arXiv:2401.00096*, 2023.
- [47] Chi Chen and Shyue Ping Ong. A universal graph deep learning interatomic potential for the periodic table. *Nature Computational Science*, 2(11):718–728, 2022.
- [48] Albert Musaelian, Simon Batzner, Anders Johansson, Lixin Sun, Cameron J Owen, Mordechai Kornbluth, and Boris Kozinsky. Learning local equivariant representations for large-scale atomistic dynamics. *Nature Communications*, 14(1):579, 2023.
- [49] Brandon M Wood, Misko Dzamba, Xiang Fu, Meng Gao, Muhammed Shuaibi, Luis Barroso-Luque, Kareem Abdelmaqsoud, Vahe Gharakhanyan, John R Kitchin, Daniel S Levine, et al. Uma: A family of universal models for atoms. *arXiv preprint arXiv:2506.23971*, 2025.
- [50] Bowen Deng, Yunyeong Choi, Peichen Zhong, Janosh Riebesell, Shashwat Anand, Zhuohan Li, KyuJung Jun, Kristin A Persson, and Gerbrand Ceder. Overcoming systematic softening in universal machine learning interatomic potentials by fine-tuning. *arXiv preprint arXiv:2405.07105*, 2024.

A Appendix

A.1 Selected Structures

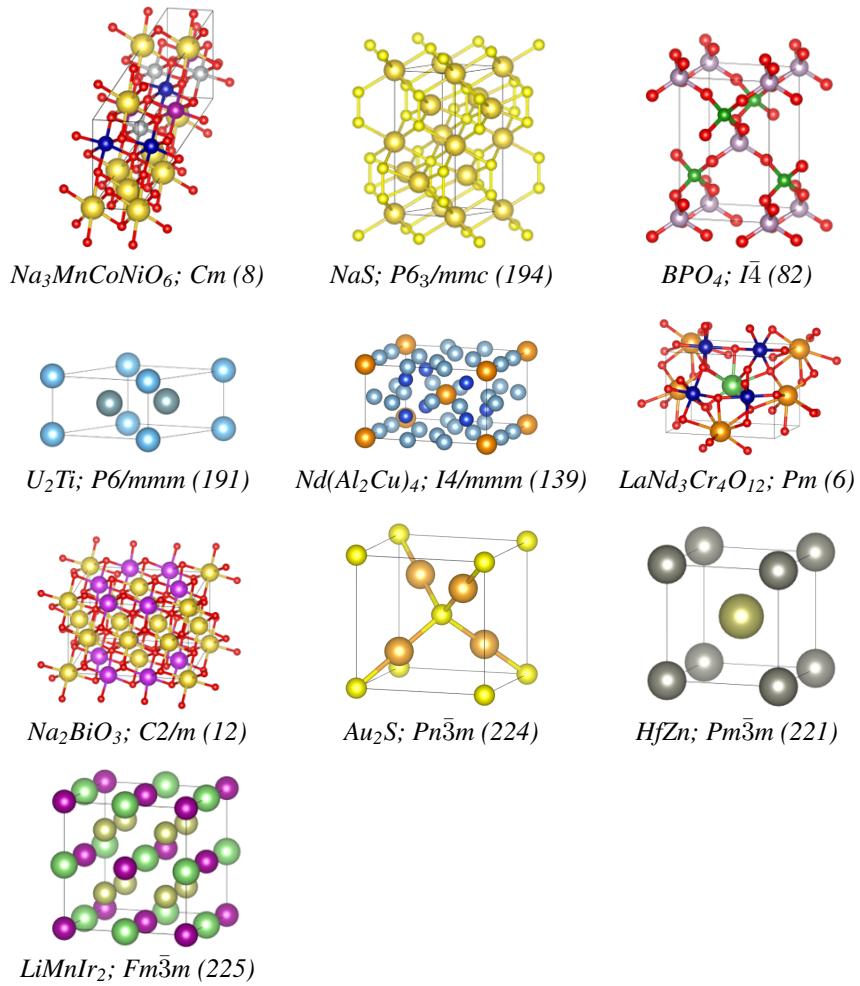


Figure 4: **Atomic structures used in this study.** Each structure is labeled with its chemical formula, space group symbol, and space group number.

A.2 XRD Representation

We follow Riesel et al. [19] and compute diffraction patterns from the structure factor contributions of each atomic site, following Bragg's law and the kinematic scattering model. Lattice parameters are converted to a real-space cell, the reciprocal lattice is derived, and all allowed Miller indices within the maximum scattering vector are generated. For each (hkl) , reciprocal distances and diffraction angles are calculated, elemental scattering factors are retrieved and weighted by site occupancies, and intensities are obtained by squaring the modulus of the summed structure factor with a Lorentz–polarization correction.

We calculate XRD peak profiles using the Pseudo-Voigt approximation, which models the peak shape as a linear combination of Gaussian and Lorentzian components:

$$pV(x) = \eta G(x) + (1 - \eta) L(x)$$

where $G(x)$ is the Gaussian function, $L(x)$ is the Lorentzian function, and $\eta \in [0, 1]$ is the mixing parameter controlling the relative contributions.

Gaussian and Lorentzian peak shapes. For a peak centered at $2\theta_0$, the Gaussian and Lorentzian components are given by:

$$G(x) = \exp \left[-\frac{4 \ln 2 (x - 2\theta_0)^2}{H_G^2} \right],$$

$$L(x) = \frac{1}{1 + \frac{4(x - 2\theta_0)^2}{H_L^2}},$$

where H_G and H_L are the full widths at half maximum (FWHM) for the Gaussian and Lorentzian profiles, respectively.

Caglioti parameters. In practice, peak broadening in XRD is described by the Caglioti relation:

$$H^2(2\theta) = U \tan^2 \theta + V \tan \theta + W,$$

where U , V , and W are the Caglioti parameters. This equation gives the squared FWHM as a function of diffraction angle, and is applied separately for the Gaussian and Lorentzian widths, i.e., $H_G(2\theta)$ and $H_L(2\theta)$. The parameters account for instrumental and sample-dependent broadening effects.

Final pattern representation. We compute the total XRD pattern by summing $pV(x)$ contributions from all Bragg reflections over $2\theta \in [0^\circ, 90^\circ]$, and then discretize the intensity into bins of width 0.01° . We adopt Caglioti parameters $U = 0.1$, $V = 0.01$, $W = 0.1$ and $\eta = 0.1$. This produces a fixed-length xrd vector, \mathbf{x} , of size 9000 for each structure.

A.3 Symmetry Projectors by Crystal Family

For each crystal family, the symmetry projector \mathcal{P} maps the given lattice parameters $(a, b, c, \alpha, \beta, \gamma)$ to the symmetrized parameters consistent with the family:

$$\begin{aligned} \mathcal{P}_{\text{cubic}}(a, b, c, \alpha, \beta, \gamma) &= (\bar{a}, \bar{a}, \bar{a}, 90^\circ, 90^\circ, 90^\circ), & \bar{a} &= \frac{a + b + c}{3} \\ \mathcal{P}_{\text{hexagonal}}(a, b, c, \alpha, \beta, \gamma) &= (\bar{a}, \bar{a}, c, 90^\circ, 90^\circ, 120^\circ), & \bar{a} &= \frac{a + b}{2} \\ \mathcal{P}_{\text{tetragonal}}(a, b, c, \alpha, \beta, \gamma) &= (\bar{a}, \bar{a}, c, 90^\circ, 90^\circ, 90^\circ), & \bar{a} &= \frac{a + b}{2} \\ \mathcal{P}_{\text{orthorhombic}}(a, b, c, \alpha, \beta, \gamma) &= (a, b, c, 90^\circ, 90^\circ, 90^\circ) \\ \mathcal{P}_{\text{monoclinic}}(a, b, c, \alpha, \beta, \gamma) &= (a, b, c, 90^\circ, \beta, 90^\circ) \\ \mathcal{P}_{\text{triclinic}}(a, b, c, \alpha, \beta, \gamma) &= (a, b, c, \alpha, \beta, \gamma) \end{aligned}$$

A.4 2D Landscape of XRD Cosine Similarity Loss

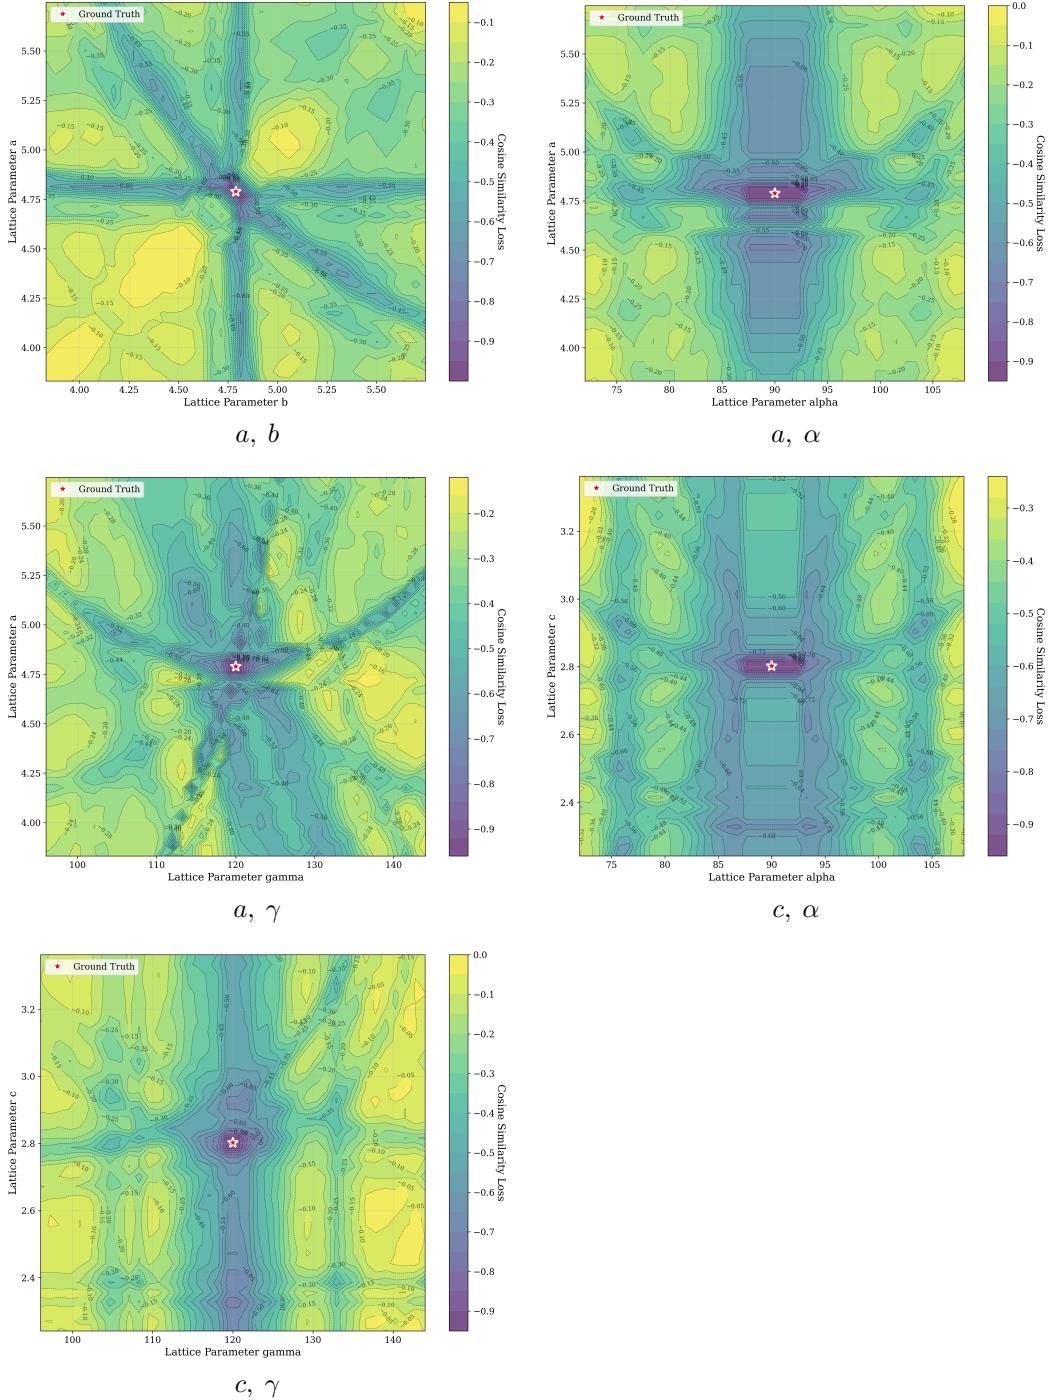


Figure 5: All 2D landscape slices of XRD cosine similarity of U_2Ti . Each slice is labeled with the distorted parameters. Due to the hexagonal symmetry of the structure, some slices are redundant and thus obscured.

A.5 Lattice and XRD Agreement

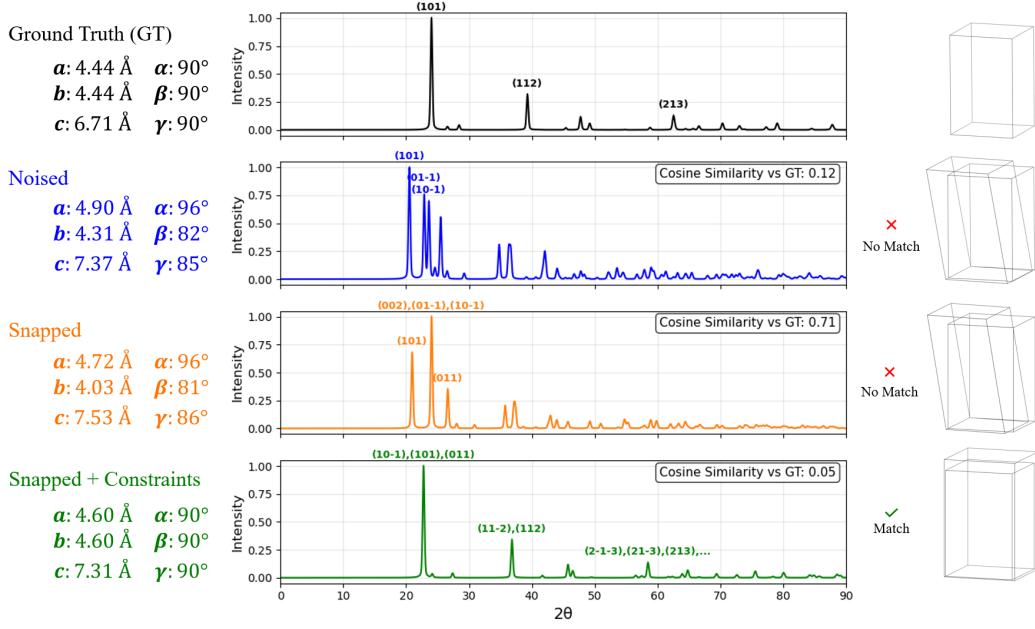


Figure 6: **Lattice and XRD patterns of BPO₄.** Each row shows the lattice parameters, corresponding XRD pattern, and unit cell relative to the ground truth. From top to bottom: ground truth; distorted lattice structure with 0.1 noise level; result of XRD-based GD optimization without constraints; and result of XRD-based GD optimization with symmetry-based constraints. For each case, the cosine similarity to the ground truth pattern and the structure match status are reported.

A.6 Comparison to Energy Optimization

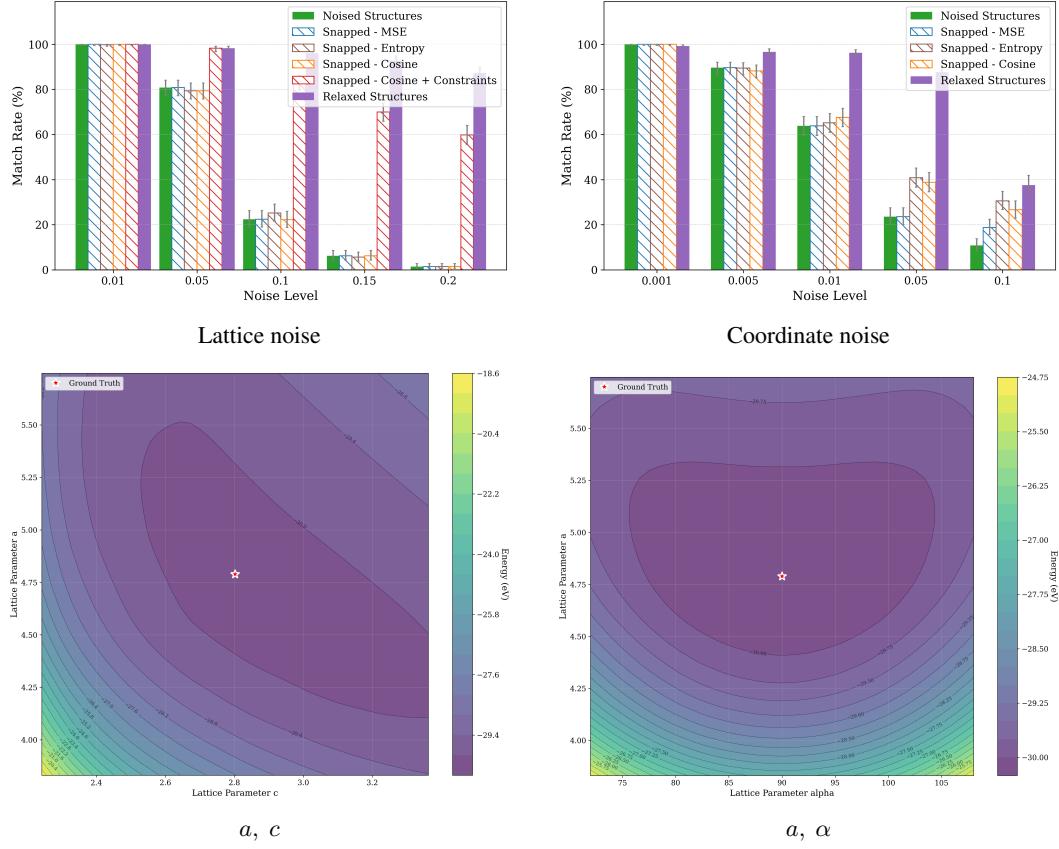


Figure 7: **Comparing XRD-based optimization with energy relaxation.** (Top) Match rates from StructureMatcher with ($l_{tol} = 0.1$, $stol = 0.2$, $angle_tol = 5^\circ$) under random lattice and coordinate perturbations. Snapped bars are the same as presented in Figure 1. Energy-based optimization consistently recovers the correct phase, except for high levels of coordinate noise, whereas XRD-based optimization struggles. (Bottom) 2D landscape of CHGNet [45] predicted energy as a function of lattice parameters of U_2Ti . Along lattice vectors, the energy landscape is smooth and locally convex.