

LEMAT-GENBENCH: A Unified Evaluation Framework for Crystal Generative Models

Siddharth Betala^{1,*}, Samuel P. Gleason¹, Ali Ramlaoui¹, Andy Xu², Georgia Channing³, Daniel Levy⁴, Clementine Fourier³, Nikita Kazeev⁵, Chaitanya K. Joshi⁶, Sekou-Oumar Kaba⁴, Felix Therrien⁴, Alex Hernandez-Garcia⁴, Rocío Mercado⁷, N. M. Anoop Krishnan⁸, Alexandre Duval^{1,*}

¹Entalpic, France, ²Harvey Mudd College, USA, ³Hugging Face, France, ⁴Mila, Canada, ⁵National University of Singapore, Singapore, ⁶University of Cambridge, UK, ⁷Chalmers University of Technology, Sweden, ⁸Indian Institute of Technology Delhi, India

*Corresponding authors

Generative machine learning (ML) models hold great promise for accelerating materials discovery through the inverse design of inorganic crystals, enabling an unprecedented exploration of chemical space. Yet, the lack of standardized evaluation frameworks makes it challenging to evaluate, compare, and further develop these ML models meaningfully. In this work, we introduce LEMAT-GENBENCH, a unified benchmark for generative models of crystalline materials, supported by a set of evaluation metrics designed to better inform model development and downstream applications. We release both an open-source evaluation suite and a public leaderboard on Hugging Face, and benchmark 12 recent generative models. Results reveal that an increase in stability leads to a decrease in novelty and diversity on average, with no model excelling across all dimensions. Altogether, LEMAT-GENBENCH establishes a reproducible and extensible foundation for fair model comparison and aims to guide the development of more reliable, discovery-oriented generative models for crystalline materials.

Code: <https://github.com/LeMaterial/lemat-genbench>

Leaderboard: <https://huggingface.co/spaces/LeMaterial/LeMat-GenBench>

Correspondence: siddharth.betala-ext@entalpic.ai, alexandre.duval@entalpic.ai

ENTALPIC

1 Introduction

Discovery of inorganic crystalline materials has traditionally relied on an Edisonian cycle of expert intuition and experimental validation [Schmidt et al., 2019], occasionally aided by quantum simulations such as density functional theory (DFT) [Kohn et al., 1996, Sholl and Steckel, 2009]. While such quantum simulations provide valuable insights into structure and stability, they remain computationally expensive and require a fully specified atomic configuration (crystal lattices^{†1} and atomic positions), information that is rarely available when proposing novel materials. Machine learning (ML) models [Deringer et al., 2019, Unke et al., 2021], particularly those based on geometric graph neural networks [Duval et al., 2023], offer fast proxies for DFT and enable scalable evaluation of candidate structures. Yet, similar to DFT, these models are limited to predicting the properties (such as energy and forces) for given atomic structure and are not designed to explore or propose new crystal geometries. This has motivated a growing wave of generative ML models for materials discovery [Mila AI4Science et al., 2023, Levy et al., 2025, Kazeev et al., 2025a, Zeni et al., 2025].

These models, ranging from variational autoencoders (VAEs) [Kingma and Welling, 2013] and diffusion models [Song et al., 2021] to GFlowNets [Bengio et al., 2023] and large language models (LLMs) [Brown et al., 2020], aim to learn the distribution of valid crystal structures and sample from it, guided by target properties when applicable. This *inverse design* paradigm promises to unlock previously inaccessible regions of chemical space and accelerate the discovery of practically relevant materials.

^{†1}Terms marked with † are defined in Box 2 of the Supplementary Material.




Figure 1 LeMat-GenBench pipeline from raw model outputs to comprehensive evaluation. The framework begins by filtering generated crystals through rigorous validity checks. Valid structures are enriched with structural fingerprints, crystallographic descriptors, and MLIP-based energetic properties. Then, LeMat-GenBench computes a unified set of metrics capturing stability, novelty, uniqueness, diversity, distributional alignment, practical synthesizability considerations, and model efficiency. The resulting metric suite provides a standardized, leaderboard-ready assessment of generative models for inorganic crystal design.

The rapid development of generative models for crystal structures has revealed a critical challenge: the absence of standardized evaluation protocols. Studies vary widely in how they assess stability, define novelty, or validate structures. For instance, researchers use different reference datasets, fingerprinting methods, energy estimators [Batatia et al., 2023, Unke et al., 2021], relaxation procedures, energy thresholds, and more, making the direct comparison of the performance of these models challenging. Without shared benchmarks designed for the fair comparison of generative models, it remains difficult to disentangle genuine model improvements from differences in evaluation design. While standardized benchmarks for predictive modeling of materials, such as Matbench [Dunn et al., 2020], have enabled progress in property prediction by providing systematic model comparison, no analogous established framework exists for the evaluation of generative models of crystal structures. This gap makes progress difficult to quantify, and leaves the community without shared reference for scientific advancement in data-driven materials design.

To this end, we introduce LEMAT-GENBENCH, a benchmarking framework aimed at standardizing the evaluation of generative models for inorganic crystal structures with the following key aspects:

- **Standardized evaluation protocol:** unified metrics suite centered on (Meta)Stable, Unique, Novel ((M.)S.U.N.) rate alongside validity, diversity, and efficiency;
- **Principled evaluation methodology:** carefully designed choices for stability estimation, validity filtering, and reference dataset construction that improve reliability and reduce common sources of error;
- **Open-source toolkit:** a public Python suite for reproducible metric computation and model evaluation;
- **Comprehensive benchmark analysis:** evaluation of 12 contemporary models, revealing clear trade-offs across metrics made available through a leaderboard to establish performances and provide a fillip to the development of better models.

By establishing shared protocols and rigorous evaluation standards, LEMAT-GENBENCH aims to enable more systematic, fair, and transparent model comparisons. The framework is not set in stone; it is designed to evolve with methodological advancements while supporting both research and practical applications in AI-driven materials.

2 Related Works

The field of generative modeling for inorganic crystals has rapidly grown, with an increasing diversity of proposed architectures, training strategies, representations, and target outputs. As the challenge of materials discovery is addressed by an expanding array of methodologies, this heterogeneity complicates the task of fair comparison. In what follows, we provide an overview of current modeling approaches (Section 2.1), followed by an analysis of existing evaluation practices and their limitations (Section 2.2). Together, these point to the need for standardized benchmarks and motivate the framework we introduce in this work.

2.1 Generative Models Overview

Generative modeling for inorganic crystals has emerged as a promising strategy for inverse design, enabling the proposal of candidate structures with desired stability, symmetry[†], or functional properties. Most of these ML models are trained on large datasets of relaxed crystal structures to generate new candidates, either by sampling from a learned distribution of valid crystals, or conditioned on specific target properties. Over the past few years, an extensive range of architectural paradigms have been explored. We briefly discuss the prominent families of models below. A more comprehensive survey of the different approaches to generative modeling crystals can be found in Section A, with a taxonomy presented in Figure 6.

Among the earliest approaches were latent-variable models. **Variational AutoEncoders** (VAEs) [Noh et al., 2019, Hoffmann et al., 2019, Court et al., 2020] encode crystal representations into a continuous latent space to enable interpolation and sampling, but struggle with decoding to valid atomic configurations [Zhao et al., 2023]. **GAN**-based methods, on the other hand, attempt to generate crystal representations via adversarial training [Kim et al., 2020]. Despite being widely used years ago, their use has been drastically reduced due to training instabilities, limited diversity, and poor domain adaptation to 3D periodic systems [Zhao et al., 2021].

Diffusion models have become the dominant paradigm, learning to gradually transform Gaussian noise into 3D representations of periodic atomic structures [Xie et al., 2021, Jiao et al., 2023, Zeni et al., 2025]. The diffusion process is able to incorporate geometric and physical inductive biases, resulting in physically plausible crystals. However, they suffer from slow inference times due to the increased number of steps required during the denoising process. **Flow-based** models offer a related but computationally faster alternative than diffusion, and learn velocity fields that map between a base distribution and the data distribution [Miller et al., 2024]. Recent variants of these models have incorporated explicit symmetry conditioning [Jiao et al., 2024, Levy et al., 2025, Chang et al., 2025, Puny et al., 2025], textual conditioning from pretrained language models [Das et al., 2025, Park et al., 2025, Sriram et al., 2024], and joint training on non-crystal datasets [Joshi et al., 2025].

An alternative paradigm uses **sequential generation** strategies, where the task of crystal generation is decomposed into a series of discrete actions or tokens. Autoregressive transformers have been used to learn stepwise generation paths of crystals using tokenizations based on Wyckoff positions[†] [Cao et al., 2024, Kazeev et al., 2025a]. **Large language models** (LLMs) have recently been adapted to crystal generation by tokenizing CIF[†] formats. They can be fine-tuned for specific tasks, and due to their pretraining, can flexibly leverage textual prompts [Gruver et al., 2024, Xu et al., 2025a, Mishra et al., 2024] and incorporate multimodal conditioning [Johansen et al., 2025, Moro et al., 2025]. **Reinforcement learning** (RL) [Zamaraeva et al., 2023, Govindarajan et al., 2024] and **Generative Flow Networks** (GFlowNets) [Mila AI4Science et al., 2023, Cipcigan et al., 2024, Podina et al., 2025] decompose generation into stepwise actions guided by a reward. They differ from the previously described generative models: rather than training on a dataset of known crystals and sampling from this learned distribution, these methods are instead directly trained to generate samples that optimize a reward, such as crystal stability and validity, or some target property.

Taken together, these methods reflect an increasingly diverse and dynamic modeling landscape for crystal generation. In designing these methods, researchers have chosen different priorities and tradeoffs, balancing aspects such as physical plausibility, efficiency, diversity, and the ability to condition on properties. These inherent variability leaves the evaluation and comparison of crystal generative models as a major open challenge—due to the lack of standardized tasks, metrics, and data splits.

2.2 Evaluation Metrics and Benchmarking Efforts

Evaluating generative models for crystal structures is inherently more complex than supervised tasks such as property prediction, where objective metrics such as Root Mean Square Error (RMSE) or Mean Absolute Error (MAE) are well established. In generative settings, one must assess both structural plausibility (i.e., whether the generated atomic arrangement could exist as a real material), and functional relevance, often in the absence of direct ground-truth targets. As a result, current evaluation protocols vary considerably, hindering systematic comparisons across the literature.

Most studies focus on generating structures that are *valid* and *stable*. Validity is typically defined via heuristic filters such as charge neutrality (requiring that the sum of oxidation states across all atoms equals zero) and interatomic distances (being greater than a target physical value) [Xie et al., 2021]. Some approaches also measure validity by comparing the space group number distribution of generated structures against reference datasets through distributional distance metrics [Baird et al., 2024]. However, these checks vary across works, serve different purposes, either as hard pre-filters or as indicative sanity checks, and are not always consistent with real-world data. For example, a notable number of experimentally observed structures in Materials Project [Jain et al., 2013] entries fail charge neutrality tests due to phenomena such as partial oxidation, which can be challenging to quantify (see Section 5).

Thermodynamic stability is usually evaluated through formation energy[†] (E_f), which measures the energy difference between the entire crystal and its constituent elements, or through energy above the convex hull[†] (E_{hull}), which quantifies how far a structure lies above the lowest-energy combination of competing phases [Batatia et al., 2023]. Studies differ in how they define stability thresholds (some requiring structures on the hull, others allowing up to 0.1 eV/atom above it), in their relaxation strategies[†], and in their choice of reference datasets. Besides, recent studies often favor Machine Learning Interatomic Potentials (MLIPs) to predict energies, since they scale better than DFT. However, MLIPs tend to underestimate E_f [Fu et al., 2022, Nong et al., 2025] and incorporate biases related to the distribution of their training datasets. Only a few works quantify this uncertainty or reconcile MLIP-based scores with DFT-based convex hulls.

Building upon stability, the *S.U.N.* framework (Stable, Unique, Novel) [Zeni et al., 2025] rewards models for generating thermodynamically viable structures (Stable) that are both distinct from one another (Unique) and absent from known databases (Novel). This metric has emerged as a composite diagnostic, but lacks a consensus implementation. Uniqueness and novelty rely on structure-based matching, using tools such as **StructureMatcher** [Ong et al., 2013] or **BAWL** [Siron et al., 2025] which compare atomic arrangements under symmetry operations to determine structure equivalence, whose results depend on the reference set, tolerance thresholds, and whether novelty is computed over all samples or only the stable ones. Most works rely on a single scheme without systematically comparing alternatives or handling disorder (e.g., partial occupancies or mixed species) [Siron et al., 2025] affecting the relevance of S.U.N. scores.

Additional metrics like diversity and distribution similarity are frequently used. Distribution similarity, often measured by Maximum Mean Discrepancy (MMD)² [Gretton et al., 2012] or Earth Mover’s Distance (EMD)³ [Rubner et al., 2000], often rewards memorization rather than novelty, misaligned with discovery goals. On the other hand, diversity lacks a unified definition, implementation, and visualization across studies. Lastly, efficiency metrics like training time, memory usage, and inference cost are under-reported, though increasingly important as generation receives more attention and model sizes grow.

To address this fragmentation, we introduce LEMAT-GENBENCH: a unified, extensible benchmark that proposes a list of standardized metrics and tasks, enabling reproducible model comparison through public tools and a live leaderboard.

²MMD compares whether samples come from the same distribution by comparing their averaged representations

³EMD computes the minimum displacement, under a chosen metric, needed to transform one distribution into another)

3 Benchmark Methodology

3.1 Evaluation Tasks

Crystal generative modeling spans a range of practical use cases, from learning broad structural priors to targeting application-specific properties. To reflect these different levels of supervision, we conceptually distinguish three evaluation scenarios: (i) *unconditional generation*, which assesses a model’s ability to sample realistic crystalline materials without guidance; (ii) *conditional generation*, where models are steered toward property-aware design using inexpensive oracle⁴ evaluations; and (iii) *conditional generation under limited budget*, which reflects settings where each oracle call (e.g., DFT or lab experiment) is costly and sample efficiency becomes paramount. We consider that these three scenarios form the conceptual landscape in which generative models for materials discovery operate.

In this work, we focus exclusively on the unconditional generation setting. This choice reflects both the maturity of unconditional modeling techniques and the need for standardized metrics to evaluate general-purpose crystal generators before they are deployed or adapted for downstream, property-driven tasks. Establishing rigorous and comparable unconditional metrics is a prerequisite to meaningfully assessing conditional or constrained discovery workflows. We therefore introduce a unified evaluation suite centered on validity, stability, novelty, uniqueness, and diversity, along with standardized reference datasets and MLIP-based energy evaluation protocols. The two conditional settings remain important for real-world discovery, which will be pursued as future extensions of LEMAT-GENBENCH.

3.2 Unconditional Generation Metrics

To evaluate unconditional generation, a representative number of structures (e.g., 2,500) are sampled from the trained model and evaluated using the following metrics, whose formal definitions are provided in [Section B](#).

Validity. We define a new validity metric, building on [Xie et al. \[2021\]](#). It serves as a pre-filtering step to exclude malformed or nonphysical structures, catching common failure modes and reducing computational overhead. These checks include CIF[†] readability, minimum interatomic distances ($> 0.7 \text{ \AA}$), and density bounds (mass density of $0.01\text{--}25 \text{ g/cm}^3$ and atomic density of $10^{-5}\text{--}0.5 \text{ atoms/\AA}^3$). We also check for reasonable lattice parameters (a, b, c within $1\text{--}100 \text{ \AA}$ and $0 < \alpha, \beta, \gamma < 180$ degrees), a crystal symmetry adhering to a space group determinable by the SpacegroupAnalyzer module in `pymatgen` [Jain et al. \[2013\]](#), and charge neutrality (via oxidation state plausibility analysis, with a tolerance threshold). All results are aggregated into a single validity score: the percentage of structures passing the validity test. Failing any of the above metrics results in an invalid label, and these structures are discarded. While submitters may pre-filter their structures, the ensuing standardized criteria ensure consistency across evaluations. An analysis of our proposed validity metric can be found in [Section 5](#).

Stability. Thermodynamic stability is a key proxy for real-world material feasibility. While formation energy has been widely used to assess stability, it introduces notable biases, such as favoring strongly bonded, electronegative compositions, and is therefore an unreliable proxy for stability. We thus utilize *energy above the convex hull* and adopt disjoint categories: structures with $E_{\text{hull}} \leq 0$ are *stable*, those with $0 < E_{\text{hull}} \leq 0.1 \text{ eV/atom}$ are *metastable*[†]. We use the term *(meta)stable* to refer to the union of both these sets, aligning with the inclusive threshold common in ML workflows.

A key methodological detail is how the convex hull is constructed. Mixing MLIP-predicted total energies with a DFT-based convex hull leads to systematic inconsistencies because MLIPs and DFT operate on different energy references depending on the MLIP’s training dataset. As demonstrated empirically (see [Section C.3](#)), such mixed-reference hulls can degrade stability prediction accuracy when de-referencing energies to make them comparable. For this reason, LEMAT-GENBENCH adopts a self-consistent MLIP-based convex hull, where each MLIP both evaluates candidate structures and defines the reference convex hull. This approach cancels potential-specific errors and yields more reliable E_{hull} estimates, particularly under strict stability thresholds.

⁴An *oracle* is a function that returns a material property score. It can be an ML model, a DFT simulation, a lab experiment or any other evaluation method. Oracles vary in computational cost and accuracy.

The convex hull is constructed from LeMat-Bulk [Siron et al., 2025], which contains approximately 5 million structures and provides the broadest coverage of competing phases relaxed under consistent DFT settings. We choose LeMat-Bulk as the reference because its breadth yields a tighter and a more realistic convex hull for stability assessment. This better reflects the needs of end-users seeking genuinely novel materials, rather than reproducing a narrow training distribution. The rationale for choosing LeMat-Bulk over datasets like MP-20 is expanded in Section 5. To enable scalable and open evaluation, we compute E_{hull} using an ensemble of MLIPs: MACE-MP [Batatia et al., 2023], UMA [Wood et al., 2025], and Orb-v3 [Rhodes et al., 2025]. For each MLIP, we construct its own convex hull and report the mean and standard deviation of E_{hull} across models, providing robustness and uncertainty estimates. To avoid expensive on-the-fly computation, we pre-compute MLIP energies for all LeMat-Bulk structures; at evaluation time, we simply filter to the relevant compositional subspace and construct the phase diagram for each generated candidate.

We additionally assess structural quality via a relaxation check: each structure is relaxed with every MLIP in the ensemble, and we compute the root mean square deviation (RMSD) between the initial and relaxed atomic positions (averaged over each MLIP). Low RMSD indicates that the submitted structure already lies near a local energy minimum. We report both *pre-relaxation* scores (computed on submitted structures) and *post-relaxation* scores (computed after MLIP relaxation); full results for both settings appear in Section C. For the public leaderboard, we favor pre-relaxation evaluation: users are encouraged to submit structures that have already been relaxed as part of their generation pipeline, ensuring that stability metrics reflect the model’s actual outputs rather than post-hoc refinement, therefore improving comparability across submissions.

Novelty measures the fraction of generated crystals not found in a reference dataset, using structural fingerprints, to identify previously unseen materials. Again, we use LeMat-Bulk [Siron et al., 2025] as the reference database for novelty evaluation: the same 5 million structures that define competing phases for the convex hull also define what counts as “known” for novelty assessment. To assess if two 3D structures point to the same material, we apply the MatterGen-adapted **StructureMatcher** [Zeni et al., 2025], partly because it handles symmetry, disorder, and known edge cases reliably. To ensure scalability, we restrict comparisons to structures of matching composition. For workflows requiring comparison across large numbers of model outputs, we also implement the Short-BAWL fingerprint [Siron et al., 2025] as a computationally efficient alternative within the codebase. Note that novelty should be treated with caution. Structural difference does not always imply functional difference, and the boundary between truly novel and incrementally different materials is inherently fuzzy. Still, standardizing fingerprint methods, reference datasets, and thresholds is critical for enabling meaningful comparisons and tracking progress.

Uniqueness quantifies non-redundancy among generated structures using the same fingerprinting approach as novelty, yielding the percentage of unique structures. Importantly, we apply validity, uniqueness and novelty sequentially, i.e., we also compute novelty on generated structures that are valid and unique, which offers different information compared to doing it on the whole set of structures.

(M.)S.U.N. rate. We aggregate core evaluation criteria into a single standardized metric following the S.U.N. framework introduced by MatterGen [Zeni et al., 2025]: the S.U.N. rate measures the fraction of generated structures that are stable, unique, and novel, computed sequentially from the submitted set in a funnel-like manner. This serves as our primary benchmark for unconditional generation in LEMAT-GENBENCH, with each component precisely defined using fixed thresholds, reference datasets, and fingerprinting methods. To account for synthesizable but metastable materials, we also report the M.S.U.N. rate, relaxing stability to $0 < E_{\text{hull}} \leq 0.1$ eV/atom. (M.)S.U.N. combines both and thus sets a practical upper bound on generative performance and provides a consistent, interpretable measure to compare models.

Distribution metrics capture the spread of generated samples across lattice parameters, space groups, elemental compositions, and structural descriptors. We compute the per-feature Shannon entropy and provide options for visualization tools in the LEMAT-GENBENCH codebase.

Diversity metrics compare the distribution of generated structures against reference datasets across key crystallographic and compositional attributes. We elaborate further on the diversity and distribution-based metrics offered as part of LEMAT-GENBENCH in Section C.

Model efficiency. Beyond output quality, we also report efficiency metrics related to training and inference. Specifically, authors must provide: (i) training compute in FLOPs and time (CPU/GPU days), (ii) inference

time to generate 2500 structures on a reference CPU/GPU (e.g., Nvidia A100)⁵, and (iii) peak memory usage during inference, together with number of model parameters. Reporting these values will inform model cost and scalability, and support analysis of tradeoffs between performance (e.g., (M.)S.U.N. rate) and deployment feasibility.

4 Towards an Open Benchmark Framework

4.1 Leaderboard Implementation

To converge towards a community-wide benchmarking framework, we adopt several concrete steps. Specifically, we provide a fully open-source implementation of the proposed evaluation metrics for unconditional evaluation, which can be used and updated by the community as the field evolves. Our contributions include: (i) LeMat-Bulk as the reference dataset for (M.)S.U.N, (ii) an ensemble of MLIPs for robust formation energy prediction, with uncertainty quantification, (iii) energy above hull calculations between an MLIP prediction and its corresponding convex hull, instead of a MLIP vs DFT-hull; (iv) enhanced validity criteria, (v) standardized diversity scores, (vi) resource efficiency reporting, (vii) pre-relaxation and post-relaxation metrics, (viii) RMSD check, (ix) a [public leaderboard on Hugging Face](#) to further facilitate fair model comparisons.

For the leaderboard, the submission process is as follows: (i) authors submit 2500 generated crystal structures; (ii) authors may also submit their packaged model, granting a compliance badge (optional); (iii) evaluation metrics are computed using our open reference implementation; (iv) results are displayed with multiple views to support comparison across tasks and generation scenarios. The code is available at <https://github.com/LeMaterial/lemat-genbench>.

This evaluation framework aims to provide clear, fair, and rigorous standards for evaluating generative models in crystal generation, while remaining flexible to evolving research needs and application contexts. Going further, we aim to extend this benchmark to the conditional generation tasks, to further bridge the gap between computational prediction and experimental realization. This infrastructure becomes particularly important as the field moves toward closed-loop discovery pipelines that integrate computational prediction, experimental validation, and synthesis planning.

4.2 Benchmarking Results and Discussion

We evaluate 12 generative models for crystalline materials on LEMAT-GENBENCH. The models considered are: ADiT [Joshi et al., 2025], Crystal-GFN [AI4Science et al., 2023], CrystalFormer [Cao et al., 2024], DiffCSP [Jiao et al., 2023], DiffCSP++ [Jiao et al., 2024], LLaMat2-CIF, LLaMat3-CIF [Mishra et al., 2024], MatterGen [Zeni et al., 2025], PLaID++ [Xu et al., 2025b], SymmCD [Levy et al., 2025], WyFormer [Kazeev et al., 2025a], and WyFormer-DFT [Kazeev et al., 2025a]. All models were trained on the MP-20 dataset. These span a broad range of architectures: diffusion models (DiffCSP, DiffCSP++, MatterGen, SymmCD), autoregressive transformers (CrystalFormer, WyFormer, WyFormer-DFT), LLM-based approaches (LLaMat2-CIF, LLaMat3-CIF), flow matching with VAE (ADiT), RL-guided models (PLaID++) and GFlowNets (Crystal-GFN). For each model, we evaluate 2500 generated structures obtained either directly from the authors or from public repositories [Kazeev et al., 2025a]. CrystalFormer is the only exception, for which only 1000 structures were available. A detailed listing of all data sources is provided in Table 3. The percentages of novel, unique, and stable structures are calculated over all submitted structures but only count valid structures, effectively setting an upper bound on achievable performance equal to the validity rate. This prevents inflated scores for models with low validity (e.g., novelty appearing high simply because the valid set is small) and ensures the reported values reflect the true frequency of valid-and-novel (or stable, unique, etc.) structures in the model’s full output distribution.

To provide reference points for interpreting generative model performance, we establish baselines by directly sampling structures from existing materials databases. These baselines are not intended as fair comparisons, but rather to establish upper bounds on certain metrics (validity, stability) and lower bounds on others

⁵While inference cost is typically negligible compared to the cost of downstream validation (e.g., DFT or experiments), standardized reporting of these metrics provides important context for practical usability and future integration into closed-loop discovery pipelines.

(novelty). They help contextualize what performance levels are achievable when simply reproducing known structures versus genuinely discovering new ones. We sample structures from MP-Exp (MP-20 entries marked as experimental by setting the ‘Theoretical’ indicator to false), AFLOW [Curtarolo et al., 2012], Alexandria [Schmidt et al., 2021, 2024], and OQMD [Kirklin et al., 2015, Saal et al., 2013]. For each database, we first randomly sample 10,000 structures without replacement, then draw 2,500 samples with replacement from this subset. This two-step procedure emulates the sampling behavior of generative models, which naturally produce duplicates when sampling from a learned distribution. Sampling with replacement ensures that uniqueness metrics remain meaningful—without it, uniqueness would trivially be 100%, unlike real generative models.

To ensure evaluations remain consistent, scalable, and physically meaningful, all metrics are computed through a simple sequential pipeline: we first apply a standardized validity check (CIF readability, density bounds, charge neutrality, minimum distances, symmetry consistency), and then compute stability, fingerprint-based, distributional, and diversity metrics on the valid subset of generated structures. This guarantees that all downstream quantities are computed on physically plausible samples and that heavy computations—such as MLIP relaxation or structural matching—remain tractable. Importantly, all evaluations are conducted on the submitted structures *without re-relaxation*. While this simplifies benchmarking and ensures reproducibility, it likely underestimates metrics such as S.U.N. and M.S.U.N. We encourage future submissions to include relaxed structures to better reflect thermodynamic viability. More detailed results can be found in Section 5 and Section C, where we report results of using MP-20 as reference set Section C.4 and of post-relaxation metrics (Table 9, Table 10, Table 15, Table 16).

Table 1 Core model evaluation metrics calculated using **MLIP ensemble** with **LeMat-Bulk** as the reference set. Submissions are organized into three categories separated by horizontal lines: (1) models that incorporate relaxation as part of their generation pipeline (MatterGen, PLaID++, WyFormer, WyFormer-DFT), (2) other generative models, (3) samples from real-world datasets. Best values are shown in **bold**, second-best values are underlined. Arrows indicate whether higher ($\uparrow$) or lower ($\downarrow$) values are better. The energy-based metrics are computed over the full generated sets which pass the initial validity check, and the relatively high standard deviations reflect the stochasticity in the structural generation process rather than measurement error..

Model	Valid% $\uparrow$	Unique% $\uparrow$	Novel% $\uparrow$	Energy-based			Stability		Metastability	
				E_f (Std) $\downarrow$	E_{hull} (Std) $\downarrow$	RMSD (Std) $\downarrow$	Stable% $\uparrow$	SUN% $\uparrow$	Metastable% $\uparrow$	MSUN% $\uparrow$
MatterGen	95.7	95.1	70.5	-0.70 ± 0.79	0.18 ± 0.18	0.39 ± 0.50	2.0	0.2 (6)	33.4	15.0
PLaID++	96.0	77.8	24.2	-0.50 ± 0.44	0.09 $\pm$ 0.16	0.13 $\pm$ 0.29	12.4	1.0 (26)	60.7	7.6
WyFormer	93.4	93.0	<u>66.4</u>	-0.43 ± 0.95	0.50 ± 0.51	0.81 ± 0.98	0.5	0.1 (3)	15.7	1.9
WyFormer-DFT	95.2	<u>95.0</u>	<u>66.4</u>	-0.67 ± 0.91	0.27 ± 0.36	0.42 ± 0.60	<u>3.7</u>	<u>0.4 (9)</u>	24.8	7.8
ADiT	90.6	87.8	26.0	-0.73 ± 0.93	0.33 ± 0.45	<u>0.38 $\pm$ 0.40</u>	0.4	0.0 (0)	<u>36.5</u>	1.0
Crystal-GFN	51.7	51.7	51.7	-1.30 ± 2.63	2.09 ± 2.38	1.87 ± 0.97	0.0	0.0 (0)	0.0	0.0
Crystalformer	69.9	69.4	31.8	-0.17 ± 1.48	0.70 ± 1.33	0.66 ± 1.05	1.4	0.0 (0)	28.8	3.1
DiffCSP	<u>95.7</u>	<u>94.8</u>	66.2	-0.64 ± 0.96	0.28 ± 0.56	0.59 ± 0.63	2.3	0.1 (3)	29.8	<u>8.5</u>
DiffCSP++	95.3	95.1	62.0	-0.52 ± 0.87	0.41 ± 0.44	0.69 ± 0.82	1.0	0.2 (4)	26.4	5.0
LLaMat2-CIF	84.4	81.4	30.0	-0.47 ± 1.01	0.44 ± 0.71	0.54 ± 0.71	0.7	0.1 (3)	34.7	2.1
LLaMat3-CIF	15.4	15.2	10.5	0.74 ± 2.69	1.71 ± 2.48	1.01 ± 1.10	0.1	0.0 (0)	2.1	0.2
SymmCD	73.4	73.0	47.0	-0.02 ± 1.22	0.88 ± 1.07	0.87 ± 1.09	1.4	0.1 (2)	18.6	2.4
Alexandria	93.4	93.4	0.0	-0.18 ± 0.94	0.44 ± 0.61	0.14 ± 0.32	2.3	0.0 (0)	27.4	0.0
OQMD	96.8	96.4	0.0	-0.23 ± 0.88	0.39 ± 0.46	0.14 ± 0.28	5.3	0.0 (0)	29.1	0.0
AFLOW	91.4	91.4	21.5	-0.19 ± 0.86	0.35 ± 0.42	0.12 ± 0.44	9.3	0.0 (0)	30.4	0.7
MP-Exp	98.0	85.6	1.9	-1.02 ± 0.97	0.09 ± 0.45	0.20 ± 0.35	15.7	0.3 (8)	64.1	0.6

Key takeaways. Tables 1 and 2 show that no single model dominates across all evaluation dimensions. Instead, models occupy different regions of the trade-off space between thermodynamic stability, novelty, distribution fidelity and diversity, which reinforces the need for multi-faceted benchmarking in generative materials design. Validity is often pretty high across models (typically $> 90\%$), indicating that generative approaches have mostly solved low-level structural correctness.

Among generative models, MatterGen attains the highest novelty fraction (70.5%) and the strongest M.S.U.N. rate (15.0%), indicating that it frequently discovers metastable structures that are also unique and novel relative to LeMat-Bulk. WyFormer-DFT offers a balanced trade-off profile, combining strong novelty and diversity with reasonable energy quality. Diffusion-based approaches demonstrate that continuous denoising processes can effectively explore regions beyond the training distribution. In contrast, Crystal-GFN and PLaID++ are clearly (over)optimised for formation energy and stability, respectively. The latter even shows the best stability metrics and S.U.N., but has substantially lower novelty (24.2%) and uniqueness (77.8%)




Figure 2 Spider plots comparing generative models using LeMat-Bulk as reference. Left: (M)S.U.N. metrics measuring validity, uniqueness, novelty, (meta)stability, and the combined (M)S.U.N. score. Right: Quality metrics including energy above hull (E_{hull}), structural relaxation (RMSD), diversity (average of elemental, space group, and size diversity normalized), JS divergence (Jensen-Shannon divergence measuring distributional similarity to reference). All metrics normalized to same scale where outer positions indicate better performance.

than MatterGen and WyFormer-like models, suggesting a tighter focus on well-explored regions of the MP-20 distribution. Models which incorporate relaxation in their generation pipelines (MatterGen, PLaID++, WyFormer, WyFormer-DFT) achieve systematically lower RMSD values than purely unrelaxed approaches. We examine the interplay between relaxation strategy and benchmark performance in [Section 5.3](#). Overall, these results show that generative models have substantial scope for improvement, as their S.U.N. rate never exceeds 1%. More broadly, the benchmark reveals that no single model excels across all dimensions: strong stability often comes at the cost of novelty, and high exploration capacity does not guarantee thermodynamic feasibility. This underscores the value of multi-metric evaluation in capturing the diverse trade-offs inherent to materials discovery. A visual comparison of the different models can be found in [Figure 2](#).

Dataset baselines. Because LeMat-Bulk unifies multiple major resources including Alexandria [[Schmidt et al., 2021, 2024](#)] and OQMD [[Kirklin et al., 2015, Saal et al., 2013](#)], baseline samples drawn from these databases provide interpretive anchors for generative model performance. As expected, dataset baselines deliver very relaxed structures with near-perfect validity and uniqueness, strong stability metrics but weak novelty, since these structures already appear in the LeMat-Bulk reference. Their distribution-matching scores (particularly for Alexandria, with the lowest JS and MMD) effectively define the reference behavior that generative models aim to approximate, as shown in [Table 2](#). AFLOW [[Curtarolo et al., 2012](#)] provides a complementary baseline: because it does not contribute directly to LeMat-Bulk, it attains non-zero novelty (21.5%) alongside strong stability (9.3% stable, 30.4% metastable), illustrating the performance of a large, external real dataset under our benchmark. Its relatively high Herfindahl-Hirschman Index (HHI)⁶ scores reflect a relatively higher representation of lower abundance/higher cost elements, emphasizing that real-world materials collections can also be composed of elements which are not fully supply chain optimized. The set of real synthesized materials (MP-Exp) outperforms all generative models and computational databases on SGDiv and SiteDiv, suggesting that generative models still underexplore the diversity of crystallographic environments found in nature. Together, these baselines establish that strong performance approaching dataset levels on stability indicates realistic structure generation, while strong novelty performance signals genuine exploration beyond known materials.

⁶Quantifies supply risk concentration for materials by measuring the concentration of element production sources and reserves. More details on how it is computed in [Section B](#)

Diversity metrics in Table 2 reveal additional aspects of model behavior. WyFormer achieves some of the highest elemental and site diversity (ElemDiv and SiteDiv), demonstrating that Wyckoff-based autoregressive architectures are particularly effective at exploring a wide range of compositions and symmetry environments. SymmCD and DiffCSP++ achieve strong space-group diversity (SGDiv), which is consistent with their explicit treatment of symmetry. Crystal-GFN reaches the highest diversity and the lowest HHI values among generative models, but at the cost of lower stability and validity metrics, highlighting that unconstrained chemical exploration can generate many structurally or thermodynamically implausible candidates. More broadly, we observe that models with very high diversity and novelty often incur larger JS, MMD, and FID values, confirming that aggressive exploration of chemical space naturally deviates from the reference training distribution. Additionally, higher novelty does not guarantee higher diversity. Mattergen illustrates this: novelty measures how combined structural and compositional attributes deviate from the reference dataset, not each attribute independently.

Taken together, these results show how **architectural choices** map to distinct performance profiles. LLM-based approaches with reinforcement learning (PLaID++) produce great stability and validity but tend to concentrate in a narrower region of chemical space. Diffusion-based models (DiffCSP, DiffCSP++, MatterGen) strike a balance between stability and novelty, with symmetry-aware variants improving physical validity and distributional alignment. Autoregressive Wyckoff-based transformers (WyFormer, WyFormer-DFT) and Crystal-GFN systematically enhance diversity and novelty by respectively encoding crystallographic constraints directly in the representation and leveraging the exploration property offered by GFlowNets [Bengio et al., 2023]. By exposing these trade-offs quantitatively, LEMAT-GENBENCH enables users to choose models based on their discovery priorities—whether exploration, thermodynamic plausibility, structural diversity, or supply-chain robustness.

Table 2 Model evaluation metrics calculated using **MLIP ensemble** with **LeMat-Bulk** as the reference set: distribution, diversity, and HHI. Submissions are organized into three categories separated by horizontal lines: (1) models that incorporate relaxation as part of their generation pipeline (MatterGen, PLaID++, WyFormer, WyFormer-DFT), (2) other generative models, (3) samples from real-world datasets. Best values are shown in **bold**, second-best values are underlined. Arrows indicate whether higher ($\uparrow$) or lower ($\downarrow$) values are better. Lower values are better for distribution metrics: JS (Jensen–Shannon divergence), MMD (maximum mean discrepancy), and FID (Fréchet Inception distance) and for HHI (Herfindahl–Hirschman index) metrics. Higher values indicate greater diversity across ElemDiv (elemental diversity), SGDiv (space-group diversity), SizeDiv (crystal size diversity), and SiteDiv (atomic site diversity). “Prod” and “Res” denote production- and reserve-side HHI components, respectively. These metrics are described in detail in Section 3.2.

Model	Distribution			Diversity				HHI	
	JS $\downarrow$	MMD $\downarrow$	FID $\downarrow$	ElemDiv $\uparrow$	SGDiv $\uparrow$	SizeDiv $\uparrow$	SiteDiv $\uparrow$	Prod $\downarrow$	Res $\downarrow$
MatterGen	0.44	0.006	2.13	0.64	0.14	0.24	12.78	3.52	2.72
PLaID++	0.43	0.029	2.69	0.68	0.26	0.21	6.01	5.23	3.36
WyFormer	0.22	0.006	<u>1.55</u>	<u>0.74</u>	<u>0.50</u>	0.27	<u>23.84</u>	3.55	2.73
WyFormer-DFT	0.25	0.010	1.72	0.73	0.49	0.26	23.56	3.49	2.74
ADiT	0.51	0.002	1.72	0.71	0.03	0.23	14.18	3.52	2.78
Crystal-GFN	0.55	0.121	43.17	0.75	0.29	0.32	32.06	<u>3.48</u>	2.34
Crystalformer	0.26	<u>0.003</u>	2.31	0.71	0.35	<u>0.31</u>	18.59	3.78	2.82
DiffCSP	0.45	0.008	2.25	0.73	0.14	0.24	14.42	3.29	<u>2.60</u>
DiffCSP++	0.21	0.003	1.28	0.72	0.51	0.27	20.84	3.56	2.77
LLaMat2-CIF	0.32	0.002	1.23	0.71	0.25	0.24	9.86	3.94	2.86
LLaMat3-CIF	0.34	0.005	9.48	0.70	0.10	0.29	13.15	3.83	2.81
SymmCD	<u>0.22</u>	0.006	1.65	0.73	0.49	0.27	19.45	3.54	2.74
Alexandria	0.09	0.000	1.31	0.71	0.50	0.25	13.50	4.02	3.33
OQMD	0.27	0.004	0.83	0.63	0.48	0.24	12.57	4.22	3.29
AFLOW	0.28	0.004	1.78	0.69	0.46	0.25	10.07	4.12	3.22
MP-Exp	0.33	0.021	2.90	0.72	0.71	0.28	54.79	2.78	2.27

5 Benchmark Design Rationale

The metrics introduced in Section 3.2 are not purely intrinsic properties of a model: they depend on concrete methodological choices in the evaluation pipeline. For example, stricter validity filters change how many structures proceed to stability checks, the way we construct the convex hull shifts stability thresholds, and

the choice of reference dataset directly affects both novelty and stability via the phase diagram. Therefore, in this section we motivate and empirically justify four key design decisions in LEMAT-GENBENCH: the validity criteria, the self-consistent MLIP-based hull construction, the pre- vs. post-relaxation protocol, and the choice of reference dataset.

5.1 Validity




Figure 3 Validity comparison between SMACT and our proposed method. Five random samples of 1,000 structures from LeMat-Bulk were evaluated using both methods. **Left:** Overall validity rates with standard deviation across seeds. **Right:** Agreement breakdown. Our method recovers 12.6% of structures that SMACT rejects, primarily f-block and metalloid containing structures and elements with multiple oxidation states, while disagreeing on only 1.0% in the opposite direction. Seed-level results in Figure 8.

Existing validity checks often reject a large fraction of experimentally confirmed structures, penalizing generative models for producing chemically plausible outputs and distorting benchmark comparisons. This is especially true for SMACT [Davies et al., 2019], which only marks $\sim 83\%$ of LeMat-Bulk structures [Siron et al., 2025] as valid, with most invalid labels arising from strict charge-neutrality assumptions, fixed oxidation state rules, and metallic structures improperly assessed using oxidation state assignments. Such rigidity causes systematic false negatives for particularly mixed-valence structures, alloys, and structures containing f-block elements.

Our hierarchical validity metric addresses these issues by incorporating empirical oxidation-state frequencies from LeMat-Bulk and ICSD and by introducing a flexible oxidation state assignment protocol for structures failing these strict checks. This enables more flexible treatment of metallic structures and structures containing ambiguous and uncommon oxidation states. As shown in Figure 3 and Figure 8, this approach increases validity from 82.7% (SMACT) to 94.3% while maintaining a low $\sim 1\%$ false-positive rate. It correctly recovers many structures, such as rare-earth oxides and metalloid containing structures, that SMACT incorrectly rejects.

Because this metric relies on empirical priors, it may slightly favor well-documented elements, but the low false-positive rate suggests this bias is limited in practice. Importantly, many valid structures may still contain scarce, toxic, or potentially radioactive elements; rather than marking them as invalid, we manage these undesirable compositions downstream through the HHI index. This separation ensures that validity reflects structural and chemical soundness, while compositional preferences remain application-specific.

5.2 Self-Consistent MLIP Convex Hull

Once structures pass validity screening, stability assessment relies on computing their energy above the convex hull (E_{hull}). Since DFT computations are expensive, a widely used shortcut is to predict candidate energies with an MLIP while comparing them against a DFT-derived hull from a reference dataset. This




Figure 4 F1-score for stability prediction on LeMat-Bulk. Three MLIPs (`orb-v3`, `uma-s1p1`, `mace-mp`) are evaluated at different E_{hull} thresholds. **(a)** MLIP energies compared against a DFT-constructed hull. **(b)** Self-consistent approach where each MLIP defines its own hull. The self-consistent method yields consistently higher F1-scores, particularly at strict thresholds.

mixed-reference setup introduces systematic inconsistencies: MLIPs and DFT operate on different energy baselines determined by their respective training sets and reference choices. Small shifts, even a few tens of meV/atom, can flip stability labels. This is problematic because E_{hull} is often treated as a hard decision boundary when evaluating the quality of structures coming out of generative models.

To eliminate these inconsistencies, LEMAT-GENBENCH adopts a self-consistent approach: each MLIP both predicts energies and defines its own convex hull. To compare this approach against the mixed-reference strategy, we treat DFT stability labels from LeMat-Bulk as ground truth. A structure is considered truly stable if its DFT energy lies within a given threshold of the DFT convex hull. We then evaluate how well each MLIP recovers these labels using either (i) the DFT hull or (ii) its self-consistent hull.

As shown in Figure 4, F1-scores⁷ improve markedly under the self-consistent setup e.g., Orb-v3: 0.66 $\rightarrow$ 0.81, UMA: 0.67 $\rightarrow$ 0.83 at $E_{\text{hull}} < 0.01$ eV/atom. The gain likely stems from model-specific energy biases canceling out when the same potential defines both candidate energies and its reference hull. Additional analyses are provided in Section C.3. Overall, these results show that the self-consistent MLIP-based hull is systematically more reliable for stability assessment and should be preferred for benchmarking generative models.

5.3 Pre-Relaxation vs. Post-Relaxation

All results in Table 1 are computed on the structures exactly as generated, without applying any additional relaxation. This is crucial for fair comparison: it avoids introducing the bias of a single external relaxer and directly evaluates the geometric quality of the structures produced by each model. In practice, high-quality generative pipelines should already output relaxed or near-relaxed structures, so evaluating them in their raw form is both more meaningful and more desirable. Nevertheless, since some models explicitly integrate relaxation into their generation pipeline, while others do not, we also compute a post-relaxations metrics. In particular, we apply a uniform UMA relaxation to all generated structures (Table 9), to understand how relaxation affects evaluation.

Models that internally relax their outputs show minimal change. PLaid++ and AFLOW display the lowest pre-relaxation RMSD values (0.10 and 0.11 Å), confirming that outputs already reside near MLIP energy minima.

Models without internal relaxation change substantially after refinement. ADiT shows the strongest shift: stability increases from 0.1% to 17.7% and metastability from 35.4% to 62.6%. DiffCSP improves similarly (stability: 2.0% to 5.9%, M.S.U.N.: 8.2% to 19.4%). These gains indicate that several models produce

⁷ $F_1 = 2 \frac{\text{precision} \cdot \text{recall}}{\text{precision} + \text{recall}}$, where precision is the fraction of predicted-stable structures that are stable according to DFT and recall is the fraction of truly stable structures correctly identified.

geometrically plausible but not fully relaxed structures, and thus benefit greatly from an additional energy minimization step.

Given these large differences, we report in this paper both pre- and post-relaxation results (Table 7, 9), since the magnitude of the shift carries important diagnostic information about each model’s geometry quality and reliance on downstream refinement. This being said, the LEMAT-GENBENCH leaderboard adopts pre-relaxation metrics as the default, to reflect the true generative capability of each model and encourage pipelines that produce well-relaxed structures directly.

5.4 Reference Dataset: LeMat-Bulk vs MP-20

Choosing a reference dataset requires balancing two objectives. Model developers benefit from evaluations aligned with the training distribution, where differences isolate architectural and algorithmic contributions. End-users, however, require metrics that reflect real discovery potential, where stability, novelty, and phase competition are assessed against the broadest available set of known materials. To serve both needs, we report MP-20 results in Section C.4 but adopt LeMat-Bulk as the primary reference for the public leaderboard.

LeMat-Bulk (~5M structures) provides the comprehensive phase coverage needed for realistic evaluation. In contrast, MP-20’s limited compositional and structural variety leads to inflated novelty and stability estimates, since many “new” or “stable” structures simply fall outside its constrained reference set. This difference is substantial: novelty drops from 83.1% to 70.5% for MatterGen; stability rates are divided by $2\text{--}3\times$ (e.g., PLaID++: 31.0% to 12.4%; DiffCSP: 7.2% to 2.3%); and S.U.N. rates collapse further (e.g., DiffCSP: 4.1% to 0.1%). These shifts illustrate that smaller references can substantially overestimate a model’s practical usefulness.

As generative models scale toward deployment in real discovery settings, evaluation against a comprehensive reference set becomes essential. MP-20 remains valuable for controlled methodological benchmarking, but LeMat-Bulk provides a more reliable basis for assessing discovery-oriented performance. We anticipate expanding reference sets as the field matures and more large-scale, consistently curated materials databases become available.

6 Conclusions and Outlook

Generative models for crystalline materials have advanced rapidly, yet the absence of shared evaluation standards has made progress difficult to measure. LEMAT-GENBENCH addresses this gap by introducing the first unified, open benchmarking framework for unconditional crystal generation. It builds on standardized validity checks, self-consistent MLIP-based stability assessments, and well-defined measures of uniqueness, novelty, and diversity. By grounding stability and novelty evaluations in the large-scale LeMat-Bulk dataset, the benchmark offers a stringent and realistic assessment of a model’s capacity to propose genuinely new materials.

Our evaluation of 12 state-of-the-art generative models highlights the diversity of modeling trade-offs: diffusion and autoregressive architectures tend to explore more compositionally and structurally diverse regions of chemical space, while RL- and LLM-guided approaches produce more stable but less novel candidates. No model dominates all metrics, underscoring the need for multi-dimensional evaluation and improved generative paradigms. The LEMAT-GENBENCH leaderboard and open-source implementation provide a common reference point for future model development, reproducibility, and community-driven extensions. We see this as a first step toward closing the loop between computational generation and experimental validation.

Limitations and Future Directions. This release focuses on unconditional generation evaluation, with conditional generation benchmarks and expanded model implementations planned for future updates. Key challenges remain. Data quality is a persistent bottleneck: widely used datasets often lack compositional diversity, structural metadata, or negative examples necessary for robust generative model training. Most generative models assume idealized, defect-free crystals, overlooking critical phenomena like disorder, doping, and non-stoichiometry that shape real-world functionality. Moreover, stability assessment relies on ensembles of MLIPs, which, despite averaging, can deviate systematically from DFT, especially near stability thresholds. Finally, although we assess thermodynamic plausibility, our framework does not yet capture kinetic barriers,

synthesis feasibility, or real-world constraints. Bridging these gaps, especially toward synthesis-aware and experimentally grounded pipelines, will require tighter integration between data, modeling, and validation across disciplines. Environmental and sustainability considerations are discussed in [Section D](#).

7 Acknowledgments

We acknowledge the fruitful discussions and insights brought by Alfonso Amayuelas, Anuroop Sriram, Benjamin Kurt Miller, Edwin Fako, Hashim Piracha, Kishalay Das, Tristan Deleu, Victor Schmidt, Zekun Danny Ren, Aron Walsh and Tian Xie.

References

- Jonathan Schmidt, Mário RG Marques, Silvana Botti, and Miguel AL Marques. Recent advances and applications of machine learning in solid-state materials science. *npj Computational Materials*, 5(1):83, 2019.
- Walter Kohn, Axel D Becke, and Robert G Parr. Density functional theory of electronic structure. *The Journal of Physical Chemistry*, 100(31):12974–12980, 1996.
- David S. Sholl and Janice A. Steckel. *Density Functional Theory: A Practical Introduction*. Wiley, March 2009. ISBN 9780470447710. doi:[10.1002/9780470447710](https://doi.org/10.1002/9780470447710). URL <http://dx.doi.org/10.1002/9780470447710>.
- Volker L Deringer, Miguel A Caro, and Gábor Csányi. Machine learning interatomic potentials as emerging tools for materials science. *Advanced Materials*, 31(46):1902765, 2019.
- Oliver T Unke, Stefan Chmiela, Huziel E Sauceda, Michael Gastegger, Igor Poltavsky, Kristof T Schutt, Alexandre Tkatchenko, and Klaus-Robert Müller. Machine learning force fields. *Chemical Reviews*, 121(16):10142–10186, 2021.
- Alexandre Duval, Simon V Mathis, Chaitanya K Joshi, Victor Schmidt, Santiago Miret, Fragkiskos D Malliaros, Taco Cohen, Pietro Liò, Yoshua Bengio, and Michael Bronstein. A hitchhiker’s guide to geometric GNNs for 3D atomic systems. *arXiv preprint arXiv:2312.07511*, 2023.
- Mila AI4Science, Alex Hernandez-Garcia, Alexandre Duval, Alexandra Volokhova, Yoshua Bengio, Divya Sharma, Pierre Luc Carrier, Michał Koziarski, and Victor Schmidt. Crystal-GFN: sampling crystals with desirable properties and constraints. *AI for Accelerated Materials Design Workshop (NeurIPS)*, 2023.
- Daniel Levy, Siba Smarak Panigrahi, Sékou-Oumar Kaba, Qiang Zhu, Kin Long Kelvin Lee, Mikhail Galkin, Santiago Miret, and Siamak Ravanbakhsh. SymmCD: Symmetry-preserving crystal generation with diffusion models. In *The Thirteenth International Conference on Learning Representations*, 2025.
- Nikita Kazeev, Wei Nong, Ignat Romanov, Ruiming Zhu, Andrey E Ustyuzhanin, Shuya Yamazaki, and Kedar Hippalgaonkar. Wyckoff transformer: Generation of symmetric crystals. In *Forty-Second International Conference on Machine Learning*, 2025a. URL <https://openreview.net/forum?id=eFHfRQRjJo>.
- 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*, 639(8055):624–632, 2025.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes, 2013. URL <https://arxiv.org/abs/1312.6114>.
- Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=PXTIG12RRHS>.
- Yoshua Bengio, Salem Lahlou, Tristan Deleu, Edward J Hu, Mo Tiwari, and Emmanuel Bengio. GFlowNet foundations. *Journal of Machine Learning Research (JMLR)*, 2023.

- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33:1877–1901, 2020.
- Ilyes Batatia, Philipp Benner, Yuan Chiang, Alin M. Elena, Dávid P. Kovács, Janosh Riebesell, Xavier R. Advincula, Mark Asta, Matthew Avaylon, William J. Baldwin, Fabian Berger, Noam Bernstein, Arghya Bhowmik, Samuel M. Blau, Vlad Cărare, James P. Darby, Sandip De, Flaviano Della Pia, Volker L. Deringer, Rokas Elijošius, Zakariya El-Machachi, Fabio Falcioni, Edwin Fako, Andrea C. Ferrari, Annalena Genreith-Schriever, Janine George, Rhys E. A. Goodall, Clare P. Grey, Petr Grigorev, Shuang Han, Will Handley, Hendrik H. Heenen, Kersti Hermansson, Christian Holm, Jad Jaafar, Stephan Hofmann, Konstantin S. Jakob, Hyunwook Jung, Venkat Kapil, Aaron D. Kaplan, Nima Karimitari, James R. Kermode, Namu Kroupa, Jolla Kullgren, Matthew C. Kuner, Domantas Kuryla, Guoda Liepuoniute, Johannes T. Margraf, Ioan-Bogdan Magdău, Angelos Michaelides, J. Harry Moore, Aakash A. Naik, Samuel P. Niblett, Sam Walton Norwood, Niamh O’Neill, Christoph Ortner, Kristin A. Persson, Karsten Reuter, Andrew S. Rosen, Lars L. Schaaf, Christoph Schran, Benjamin X. Shi, Eric Sivonxay, Tamás K. Stenczel, Viktor Svahn, Christopher Sutton, Thomas D. Swinburne, Jules Tilly, Cas van der Oord, Eszter Varga-Umbrich, Tejs Vegge, Martin Vondrák, Yangshuai Wang, William C. Witt, Fabian Zills, and Gábor Csányi. A foundation model for atomistic materials chemistry. *arXiv preprint arXiv: 2401.00096*, 2023.
- Alexander Dunn, Qi Wang, Alex Ganose, Daniel Dopp, and Anubhav Jain. Benchmarking materials property prediction methods: the matbench test set and automatminer reference algorithm. *npj Computational Materials*, 6(1):138, 2020.
- Juhwan Noh, Jaehoon Kim, Helge S. Stein, Benjamin Sanchez-Lengeling, John M. Gregoire, Alan Aspuru-Guzik, and Yousung Jung. Inverse design of solid-state materials via a continuous representation. *Matter*, 1(5):1370–1384, 2019. ISSN 2590-2385. doi:[10.1016/j.matt.2019.08.017](https://doi.org/10.1016/j.matt.2019.08.017).
- Jordan Hoffmann, Louis Maestrati, Yoshihide Sawada, Jian Tang, Jean Michel Sellier, and Yoshua Bengio. Data-driven approach to encoding and decoding 3D crystal structures, 2019. URL <http://arxiv.org/abs/1909.00949>.
- Callum J. Court, Batuhan Yildirim, Apoorv Jain, and Jacqueline M. Cole. 3D inorganic crystal structure generation and property prediction via representation learning. *Journal of Chemical Information and Modeling*, 60(10):4518–4535, 2020. ISSN 1549-9596, 1549-960X. doi:[10.1021/acs.jcim.0c00464](https://doi.org/10.1021/acs.jcim.0c00464). URL <https://pubs.acs.org/doi/10.1021/acs.jcim.0c00464>.
- Yong Zhao, Edirisuriya M. Dilanga Siriwardane, Zhenyao Wu, Nihang Fu, Mohammed Al-Fahdi, Ming Hu, and Jianjun Hu. Physics guided deep learning for generative design of crystal materials with symmetry constraints. *npj Computational Materials*, 9(1):38, 2023. ISSN 2057-3960. doi:[10.1038/s41524-023-00987-9](https://doi.org/10.1038/s41524-023-00987-9). URL <https://doi.org/10.1038/s41524-023-00987-9>.
- Sungwon Kim, Juhwan Noh, Geun Ho Gu, Alan Aspuru-Guzik, and Yousung Jung. Generative adversarial networks for crystal structure prediction. *ACS Central Science*, 6(8):1412–1420, 2020. doi:[10.1021/acscentsci.0c00426](https://doi.org/10.1021/acscentsci.0c00426). URL <https://doi.org/10.1021/acscentsci.0c00426>. PMID: 32875082.
- Yong Zhao, Mohammed Al-Fahdi, Ming Hu, Edirisuriya M. D. Siriwardane, Yuqi Song, Alireza Nasiri, and Jianjun Hu. High-throughput discovery of novel cubic crystal materials using deep generative neural networks. *Advanced Science*, 8(20):2100566, 2021. doi:<https://doi.org/10.1002/advs.202100566>. URL <https://advanced.onlinelibrary.wiley.com/doi/abs/10.1002/advs.202100566>.
- 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.
- 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.
- Benjamin Miller, Ricky Chen, Anuroop Sriram, and Brandon Wood. FlowMM: Generating materials with Riemannian flow matching. In *International Conference on Machine Learning (ICML)*. ICML, jun 2024.

- Rui Jiao, Wenbing Huang, Yu Liu, Deli Zhao, and Yang Liu. Space group constrained crystal generation. *arXiv preprint arXiv:2402.03992*, 2024.
- Rees Chang, Angela Pak, Alex Guerra, Ni Zhan, Nick Richardson, Elif Ertekin, and Ryan P Adams. Space group equivariant crystal diffusion. *arXiv preprint arXiv:2505.10994*, 2025.
- Omri Puny, Yaron Lipman, and Benjamin Kurt Miller. Space group conditional flow matching. *arXiv preprint arXiv:2509.23822*, 2025.
- Kishalay Das, Subhojyoti Khastagir, Pawan Goyal, Seung-Cheol Lee, Satadeep Bhattacharjee, and Niloy Ganguly. Periodic materials generation using text-guided joint diffusion model. In *The Thirteenth International Conference on Learning Representations*, 2025.
- Hyunsoo Park, Anthony Onwuli, and Aron Walsh. Exploration of crystal chemical space using text-guided generative artificial intelligence. *Nature Communications*, 16(1):4379, 2025.
- Anuroop Sriram, Benjamin Miller, Ricky Chen, and Brandon Wood. FlowLLM: Flow matching for material generation with large language models as base distributions. In *Neural Information Processing Systems (NeurIPS)*. NeurIPS, jan 2024.
- Chaitanya K Joshi, Xiang Fu, Yi-Lun Liao, Vahe Gharakhanyan, Benjamin Kurt Miller, Anuroop Sriram, and Zachary W Ulissi. All-atom diffusion transformers: Unified generative modelling of molecules and materials. In *International Conference on Learning Representations*, 2025.
- Zhendong Cao, Xiaoshan Luo, Jian Lv, and Lei Wang. Space group informed transformer for crystalline materials generation. *arXiv preprint arXiv:2403.15734*, 2024.
- Nate Gruver, Anuroop Sriram, Andrea Madotto, Andrew Gordon Wilson, C Lawrence Zitnick, and Zachary Ulissi. Fine-tuned language models generate stable inorganic materials as text. *arXiv preprint arXiv:2402.04379*, 2024.
- Andy Xu, Rohan Desai, Larry Wang, Gabriel Hope, and Ethan T. Ritz. PLaID: Preference aligned language model for targeted inorganic materials design. In *AI for Accelerated Materials Design Workshop (ICLR)*, 2025a. URL <https://openreview.net/forum?id=7aoP3ZeBfy>.
- Vaibhav Mishra, Somaditya Singh, Dhruv Ahlawat, Mohd Zaki, Vaibhav Bihani, Hargun Singh Grover, Biswajit Mishra, Santiago Miret, N M Anoop Krishnan, et al. Foundational large language models for materials research. *arXiv preprint arXiv:2412.09560*, 2024.
- 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.
- Viggo Moro, Charlotte Loh, Rumen Dangovski, Ali Ghorashi, Andrew Ma, Zhuo Chen, Samuel Kim, Peter Y Lu, Thomas Christensen, and Marin Soljačić. Multimodal foundation models for material property prediction and discovery. *Newton*, 2025.
- Elena Zamaraeva, Christopher M Collins, Dmytro Antypov, Vladimir V Gusev, Rahul Savani, Matthew S Dyer, George R Darling, Igor Potapov, Matthew J Rosseinsky, and Paul G Spirakis. Reinforcement learning in crystal structure prediction. *Digital Discovery*, 2(6):1831–1840, 2023.
- Prashant Govindarajan, Santiago Miret, Jarrod Rector-Brooks, Mariano Phielipp, Janarthanan Rajendran, and Sarath Chandar. Learning conditional policies for crystal design using offline reinforcement learning. *Digital Discovery*, 3(4):769–785, 2024.
- Flaviu Cipcigan, Jonathan Booth, Rodrigo Neumann Barros Ferreira, Carine Ribeiro Dos Santos, and Mathias Steiner. Discovery of novel reticular materials for carbon dioxide capture using GFlowNets. *Digital Discovery*, 2024.
- Lena Podina, Christina Humer, Alexandre Duval, Victor Schmidt, Ali Ramlaoui, Shahana Chatterjee, Yoshua Bengio, Alex Hernandez-Garcia, David Rolnick, and Félix Therrien. Catalyst GFlowNet for electrocatalyst design: A hydrogen evolution reaction case study. *arXiv preprint arXiv:2510.02142*, 2025.

- Sterling G Baird, Hasan M Sayeed, Joseph Montoya, and Taylor D Sparks. matbench-genmetrics: A Python library for benchmarking crystal structure generative models using time-based splits of Materials Project structures. *Journal of Open Source Software*, 9(97):5618, 2024.
- 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](https://doi.org/10.1063/1.4812323).
- Xiang Fu, Zhenghao Wu, Wujie Wang, Tian Xie, Sinan Ketten, Rafael Gomez-Bombarelli, and Tommi Jaakkola. Forces are not enough: Benchmark and critical evaluation for machine learning force fields with molecular simulations. *arXiv preprint arXiv:2210.07237*, 2022.
- Wei Nong, Ruiming Zhu, Zekun Ren, Martin Hoffmann Petersen, Shuya Yamazaki, Nikita Kazeev, Andrey Ustyuzhanin, Gang Wu, Shuo-Wang Yang, and Kedar Hippalgaonkar. Energy underprediction from symmetry in machine-learning interatomic potentials. *arXiv preprint arXiv:2507.15190*, 2025.
- 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.
- Martin Siron, Inel Djafar, Ali Ramlaoui, Etienne du Fayette, Amandine Rossello, Edvin Fako, Matthew McDermott, Felix Therrien, Luis Barroso-Luque, Flaviu Cipcigan, et al. LeMat-Bulk: Aggregating, and de-duplicating quantum chemistry materials databases. *arXiv preprint arXiv:2511.05178*, 2025.
- Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *The journal of machine learning research*, 13(1):723–773, 2012.
- Yossi Rubner, Carlo Tomasi, and Leonidas J Guibas. The earth mover’s distance as a metric for image retrieval. *International journal of computer vision*, 40(2):99–121, 2000.
- 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.
- Benjamin Rhodes, Sander Vandenhaute, Vaidotas Šimkus, James Gin, Jonathan Godwin, Tim Duignan, and Mark Neumann. Orb-v3: atomistic simulation at scale. *arXiv preprint arXiv:2504.06231*, 2025.
- Mila AI4Science, Alex Hernandez-Garcia, Alexandre Duval, Alexandra Volokhova, Yoshua Bengio, Divya Sharma, Pierre Luc Carrier, Yasmine Benabed, Michał Koziarski, and Victor Schmidt. Crystal-gfn: sampling crystals with desirable properties and constraints. *arXiv preprint arXiv:2310.04925*, 2023.
- Andy Xu, Rohan Desai, Larry Wang, Gabriel Hope, and Ethan Ritz. Plaid++: A preference aligned language model for targeted inorganic materials design. *arXiv preprint arXiv:2509.07150*, 2025b.
- Stefano Curtarolo, Wahyu Setyawan, Gus LW Hart, Michal Jahnatek, Roman V Chepulskii, Richard H Taylor, Shidong Wang, Junkai Xue, Kesong Yang, Ohad Levy, et al. Aflow: An automatic framework for high-throughput materials discovery. *Computational Materials Science*, 58:218–226, 2012.
- Jonathan Schmidt, Love Pettersson, Claudio Verdozzi, Silvana Botti, and Miguel AL Marques. Crystal graph attention networks for the prediction of stable materials. *Science Advances*, 7(49):eabi7948, 2021.
- Jonathan Schmidt, Tiago FT Cerqueira, Aldo H Romero, Antoine Loew, Fabian Jäger, Hai-Chen Wang, Silvana Botti, and Miguel AL Marques. Improving machine-learning models in materials science through large datasets. *Materials Today Physics*, 48:101560, 2024.
- Scott Kirklin, James E Saal, Bryce Meredig, Alex Thompson, Jeff W Doak, Muratahan Aykol, Stephan Rühl, and Chris Wolverton. The Open Quantum Materials Database (OQMD): Assessing the accuracy of DFT formation energies. *npj Computational Materials*, 1(1):1–15, 2015.

- James E Saal, Scott Kirklin, Muratahan Aykol, Bryce Meredig, and Christopher Wolverton. Materials design and discovery with high-throughput density functional theory: The Open Quantum Materials Database (OQMD). *Jom*, 65(11):1501–1509, 2013.
- Daniel W Davies, Keith T Butler, Adam J Jackson, Jonathan M Skelton, Kazuki Morita, and Aron Walsh. SMACT: Semiconducting materials by analogy and chemical theory. *Journal of Open Source Software*, 4(38):1361, 2019.
- Zekun Ren, Siyu Isaac Parker Tian, Juhwan Noh, Felipe Oviedo, Guangzong Xing, Jiali Li, Qiaohao Liang, Ruiming Zhu, Armin G. Aberle, Shijing Sun, Xiaonan Wang, Yi Liu, Qianxiao Li, Senthilnath Jayavelu, Kedar Hippalgaonkar, Yousung Jung, and Tonio Buonassisi. An invertible crystallographic representation for general inverse design of inorganic crystals with targeted properties. *Matter*, 5(1):314–335, 2022. ISSN 2590-2385. doi:[10.1016/j.matt.2021.11.032](https://doi.org/10.1016/j.matt.2021.11.032). URL <https://www.sciencedirect.com/science/article/pii/S2590238521006251>.
- Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in Neural Information Processing Systems*, 27, 2014.
- Teng Long, Nuno M Fortunato, Ingo Opahle, Yixuan Zhang, Ilias Samathrakakis, Chen Shen, Oliver Gutfleisch, and Hongbin Zhang. Constrained crystals deep convolutional generative adversarial network for the inverse design of crystal structures. *npj Computational Materials*, 7(1):66, 2021.
- Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 2256–2265, Lille, France, 07–09 Jul 2015. PMLR.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 6840–6851. Curran Associates, Inc., 2020.
- Youzhi Luo, Chengkai Liu, and Shuiwang Ji. Towards symmetry-aware generation of periodic materials. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=Jkc74vn1aZ>.
- Sherry Yang, KwangHwan Cho, Amil Merchant, Pieter Abbeel, Dale Schuurmans, Igor Mordatch, and Ekin Dogus Cubuk. Scalable diffusion for materials generation. *arXiv preprint arXiv:2311.09235*, 2023.
- Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=PqvMRDCJT9t>.
- Michael Samuel Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic interpolants. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=li7qeBbCR1t>.
- Xingchao Liu, Chengyue Gong, and qiang liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=XVjTT1nw5z>.
- Ruiqi Gao, Emiel Hoogetboom, Jonathan Heek, Valentin De Bortoli, Kevin Patrick Murphy, and Tim Salimans. Diffusion models and gaussian flow matching: Two sides of the same coin. In *The Fourth Blogpost Track at ICLR 2025*, 2025. URL <https://openreview.net/forum?id=C8Yyg9wy0s>.
- Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT Press Cambridge, 1998.
- Kai Arulkumaran, Marc Peter Deisenroth, Miles Brundage, and Anil Anthony Bharath. Deep reinforcement learning: A brief survey. *IEEE Signal Processing Magazine*, 34(6):26–38, 2017.

- Emmanuel Bengio, Moksh Jain, Maksym Korablyov, Doina Precup, and Yoshua Bengio. Flow network based generative models for non-iterative diverse candidate generation. *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- Daniil Tiapkin, Nikita Morozov, Alexey Naumov, and Dmitry Vetrov. Generative flow networks as entropy-regularized RL. *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2024.
- Tristan Deleu, Padideh Nouri, Nikolay Malkin, Doina Precup, and Yoshua Bengio. Discrete probabilistic inference as control in multi-path environments. *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2024.
- Yu Zhang, Xiusi Chen, Bowen Jin, Sheng Wang, Shuiwang Ji, Wei Wang, and Jiawei Han. A comprehensive survey of scientific large language models and their applications in scientific discovery. *arXiv preprint arXiv:2406.10833*, 2024.
- Grégoire Delétang, Anian Ruoss, Paul-Ambroise Duquenne, Elliot Catt, Tim Genewein, Christopher Mattern, Jordi Grau-Moya, Li Kevin Wenliang, Matthew Aitchison, Laurent Orseau, et al. Language modeling is compression. *arXiv preprint arXiv:2309.10668*, 2023.
- 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.
- Luis M Antunes, Keith T Butler, and Ricardo Grau-Crespo. Crystal structure generation with autoregressive large language modeling. *Nature Communications*, 15(1):1–16, 2024.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- Pierre-Paul De Breuck, Hashim A Piracha, Gian-Marco Rignanes, and Miguel AL Marques. A generative material transformer using wyckoff representation. *arXiv preprint arXiv:2501.16051*, 2025.
- Nikita Kazeev, Wei Nong, Ignat Romanov, Ruiming Zhu, Andrey E Ustyuzhanin, Shuya Yamazaki, and Kedar Hippalgaonkar. Wyckoff Transformer: Generation of Symmetric Crystals. In Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff, and Jerry Zhu, editors, *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pages 29495–29526. PMLR, 13–19 Jul 2025b. URL <https://proceedings.mlr.press/v267/kazeev25a.html>.
- Theo Hahn, Uri Shmueli, and JC Wilson Arthur. *International tables for crystallography*, volume 1. Reidel Dordrecht, 1983.
- Ruiming Zhu, Wei Nong, Shuya Yamazaki, and Kedar Hippalgaonkar. WyCryst: Wyckoff inorganic crystal generator framework. *Matter*, 7(10):3469–3488, 2024.
- 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.
- Botian Wang, Yawen Ouyang, Yaohui Li, Yiqun Wang, Haorui Cui, Jianbing Zhang, Xiaonan Wang, Wei-Ying Ma, and Hao Zhou. Moma: A modular deep learning framework for material property prediction. *arXiv preprint arXiv:2502.15483*, 2025.
- Onur Boyar, Indra Priyadarsini, Seiji Takeda, and Lisa Hamada. LLM-Fusion: A novel multimodal fusion model for accelerated material discovery. *arXiv preprint arXiv:2503.01022*, 2025.
- Nathan J Szymanski, Bernardus Rendy, Yuxing Fei, Rishi E Kumar, Tanjin He, David Milsted, Matthew J McDermott, Max Gallant, Ekin Dogus Cubuk, Amil Merchant, et al. An autonomous laboratory for the accelerated synthesis of novel materials. *Nature*, 624(7990):86–91, 2023.

- Dejan Zagorac, H Müller, S Ruehl, J Zagorac, and Silke Rehme. Recent developments in the inorganic crystal structure database: theoretical crystal structure data and related features. *Applied Crystallography*, 52(5): 918–925, 2019.
- Shyue P Ong, Lei Wang, Byoungwoo Kang, and Gerbrand Ceder. Li-Fe-P-O₂ phase diagram from first principles calculations. *Chemistry of Materials*, 20(5):1798–1807, 2008.
- Martin Siron, Inel Djafar, Lucile Ritchie, Etienne Du-Fayet, Amandine Rossello, Ali Ramlaoui, Leandro von Werra, Thomas Wolf, and Alexandre Duval. LeMat-Bulk Dataset, 2024. URL <https://huggingface.co/datasets/LeMaterial/LeMat-Bulk>.
- Dan Friedman and Adji Bouso Dieng. The Vendi score: A diversity evaluation metric for machine learning. *arXiv preprint arXiv:2210.02410*, 2022.
- Bent Fuglede and Flemming Topsøe. Jensen-Shannon divergence and Hilbert space embedding. In *International Symposium on Information Theory (ISIT)*, page 31. IEEE, 2004.
- Ilya O Tolstikhin, Bharath K Sriperumbudur, and Bernhard Schölkopf. Minimax estimation of maximum mean discrepancy with radial kernels. *Advances in Neural Information Processing Systems*, 29, 2016.
- Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. GANs trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in Neural Information Processing Systems*, 30, 2017.
- Kristina Preuer, Philipp Renz, Thomas Unterthiner, Sepp Hochreiter, and Gunter Klambauer. Fréchet ChemNet distance: a metric for generative models for molecules in drug discovery. *Journal of Chemical Information and Modeling*, 58(9):1736–1741, 2018.
- Wenhao Gao, Tianfan Fu, Jimeng Sun, and Connor Coley. Sample efficiency matters: a benchmark for practical molecular optimization. *Advances in Neural Information Processing Systems*, 35:21342–21357, 2022.
- Aria Mansouri Tehrani, Leila Ghadbeigi, Jakoah Brgoch, and Taylor D Sparks. Balancing mechanical properties and sustainability in the search for superhard materials. *Integrating Materials and Manufacturing Innovation*, 6(1):1–8, 2017.

Supplementary Material

Box 1: Key Terms in Generative Modeling

Generative Model: A machine learning model that learns a data distribution $p(\mathbf{x})$ (or a conditional distribution $p(\mathbf{x}|\mathbf{z})$ or $p(\mathbf{x}|\mathbf{c})$) and can generate new samples $\mathbf{x}' \sim p(\mathbf{x})$ that resemble the training data.

Latent Space: A lower-dimensional representation space $\mathbf{z} \in \mathbb{R}^d$ learned by models such as VAEs or GANs, where semantic attributes of the data are often encoded.

Prior Distribution: A predefined distribution (e.g., Gaussian) over the latent variables, typically denoted as $p(\mathbf{z})$, from which samples are drawn during generation.

Decoder / Generator: A neural network (often denoted $G(\mathbf{z})$) that maps latent codes $\mathbf{z}$ to data samples $\mathbf{x}$.

Reconstruction Loss: A metric used in training autoencoders and VAEs that measures how well the generated sample $\hat{\mathbf{x}}$ matches the original input $\mathbf{x}$:

$$\mathcal{L}_{\text{recon}} = \|\hat{\mathbf{x}} - \mathbf{x}\|^2 \quad \text{or} \quad -\log p(\mathbf{x}|\mathbf{z}).$$

KL Divergence: A measure of how much one probability distribution differs from another. Commonly used in VAEs to regularize the encoder:

$$\mathcal{L}_{\text{KL}} = D_{\text{KL}}(q(\mathbf{z}|\mathbf{x})\|p(\mathbf{z})).$$

Mode Collapse: A failure mode in GANs where the generator produces samples with limited diversity, collapsing to a few modes of the data distribution.

Conditional Generation: Generation of samples $\mathbf{x}$ based on specified properties or constraints $\mathbf{c}$, e.g., $p(\mathbf{x}|\mathbf{c})$, enabling property-guided design.

Inverse Design: The process of searching the input space (e.g., structure, composition) that maps to a desired target property, often using a generative model or an optimization loop in latent space.

Diffusion Models: A class of generative models that learn to reverse a stochastic diffusion process. Data $\mathbf{x}_0$ is gradually perturbed into noise via:

$$q(\mathbf{x}_t|\mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t; \sqrt{\alpha_t}\mathbf{x}_0, (1 - \alpha_t)I).$$

and a neural network is trained to denoise $\mathbf{x}_t$ to recover $\mathbf{x}_0$ through a learned reverse process $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)$.

Score-Based Models: Closely related to diffusion models, they learn the score function $\nabla_{\mathbf{x}} \log p(\mathbf{x})$ and use Langevin dynamics or ODE solvers to sample from the data distribution.

$$\log p(\mathbf{x}) = \log p(\mathbf{z}) + \sum_{k=1}^K \log \left| \det \frac{\partial f_k^{-1}}{\partial \mathbf{x}} \right|.$$

Flow Matching: A recent generative approach that avoids training score functions or simulating diffusion. It directly learns a vector field $\mathbf{v}_\theta(\mathbf{x}, t)$ that maps noise to data through an ODE:

$$\frac{d\mathbf{x}}{dt} = \mathbf{v}_\theta(\mathbf{x}, t).$$

This method can be trained via supervised learning on synthetic trajectories or velocity fields between the base and target distributions.

Box 2: Key Terms in Crystallography & Materials Science

Crystal Lattice: A crystal structure is periodic in three dimensions. This periodicity is described by the lattice, which is defined as

$$\mathbf{L} = \{ l_1 \mathbf{a}_1 + l_2 \mathbf{a}_2 + l_3 \mathbf{a}_3 \mid l_1, l_2, l_3 \in \mathbb{Z} \},$$

where $\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3$ are basis vectors of $\mathbb{R}^3$.

Unit Cell: A unit cell is the smallest unit that can be translated to define the whole lattice. In three dimensions, it is always a parallelepiped.

Lattice Parameters: A lattice is typically defined in two ways: either as a set of three basis vectors, or as a set of lattice parameters $(a, b, c, \alpha, \beta, \gamma)$, where a, b, c are the lengths of edges of the unit cell, and α, β, γ are the angles between them.

Symmetry: An object's symmetry is given by the set of geometric transformations that map the object onto itself, leaving it invariant.

Space Group: Crystals can be classified by their symmetries. They possess the translational symmetry of their crystal lattices, and they may also have the point group symmetries of rotations and reflections within a unit cell. The combination of translational and point group symmetries can yield more transformations that a crystal can be symmetric to, including screw and glide symmetries. The full set of symmetric transformations that leave a crystal invariant defines the space group of the crystal. In three dimensions, there are 230 distinct space groups.

Wyckoff Position: Applying symmetry operations to a crystal may leave some atoms unaffected; for example, a rotation about an axis leaves atoms on the axis in the same position. The set of symmetry operations that do not move a position defines that position's site symmetry. A Wyckoff position is a set of positions that all have the same site symmetries, or conjugate site symmetries. For example, all points along a mirror plane may belong to the same Wyckoff position, while a point at the origin of a unit cell may have its own. Every point in a crystal can be assigned a Wyckoff position.

Formation Energy (E_f): The formation energy of a crystal is the difference in energy between the crystal and its constituent elements. It can be calculated using: $E_f = E_{\text{total}} - \sum_i n_i \mu_i$

Energy Above Convex Hull (E_{hull}): The convex hull gives linear combinations of known phases that represent the lowest-energy mixtures of materials; if a material has an energy above the hull ($E_{\text{hull}} > 0$), it is energetically favorable for it to decompose into a combination of stable phases and is therefore thermodynamically unstable. Mathematically, this can be represented as $E_{\text{hull}} = E_{\text{total}} - E_{\text{hull}}^{\text{min}}$. For example, the convex hull of table salt, NaCl, also includes pure stable Na, pure Cl, as well as NaCl₃. However, Na₂Cl has a higher formation energy than the combination of NaCl and pure Na, so it is unstable.

Metastable: Even if a crystal is not in its lowest possible energy state, it may still be metastable, meaning that a potential energy barrier prevents it from easily transitioning to a lower-energy state. A crystal having a low energy above the convex hull while also being at an energy minimum may indicate that it is metastable. Metastable materials are still important: for example, diamond is metastable, but does not readily convert to a lower energy state under normal conditions.

CIF: A Crystallographic Information File is a string-based encoding of a crystal that includes information such as atom positions, unit cell parameters, and chemical elements.

Relaxation Trajectory: sequence of intermediate atomic configurations produced during an energy-minimization procedure (e.g., via DFT, MLIPs, or classical force fields). Starting from an initial crystal structure, the optimizer iteratively updates atomic positions and lattice parameters to reduce the system's total energy.




Figure 5 An overview of the generative AI paradigm for candidate structure generation and optimization that underpins much of the work reviewed herein.




Figure 6 Overview of generative models for materials discovery discussed in this work. **(a)** Evolution of major model architectures over time, showing the early dominance of VAEs and the recent rise of LLMs. **(b)** Treemap of target properties optimized across models; box size reflects the proportion of papers mentioning each property. Space group, composition, lattice parameters, and formation energy are the most common. **(c)** Model distribution for unconditional (left) and conditional (right) generation tasks. Most conditional models can also perform unconditional generation, but not vice versa. Methods are grouped by their two dominant architecture families (1° and 2° model classes) for simplicity, with all colors matching panel (a). As many things have been simplified for this figure, please see [Section A](#) for details on the various methods presented. Each model is annotated with its primary input data type. **Abbreviations:** LLM = large language model; VAE = variational autoencoder; RL = reinforcement learning; NL prompt = natural language prompt; PXRD = powder X-ray diffraction. “CIF prefix” typically includes composition, space group, and lattice parameters; “Crystallographic file” refers to any file encoding structure data (e.g., XYZ, PDB, CIF).

A Survey of Generative Models for Crystalline Materials

A.1 Early Deep Learning Approaches

Variational autoencoders (VAEs) [Kingma and Welling, 2013] were among the first deep generative models applied to crystal structure generation, offering a powerful framework for encoding materials in continuous latent spaces. Early efforts by iMatGen [Noh et al., 2019] employed 3D grid-based image representations of crystal structures to learn a latent space of inorganic materials, further refined by a binary classifier that distinguished stable from unstable structures based on formation energies. This approach successfully demonstrated the ability to generate novel vanadium oxide materials. Subsequent work by Hoffmann et al. [2019] advanced this methodology by combining U-Net segmentation with VAEs in a two-step approach: first generating 3D density maps via a VAE, then segmenting these maps into 3D voxel grids using a U-Net model. Building on this foundation, Court et al. [2020] introduced conditional deep feature consistent VAEs (Cond-DFC-VAE) to enable property-guided crystal generation. A limitation of these early approaches is the use of voxel grid representations that are difficult to decode to physically meaningful atom coordinates and types. A significant advancement came from Ren et al. [2022], who developed a generalized invertible representation encoding both structural information (atom coordinates and types) and lattice parameters through reciprocal space transformations in the FTCP framework. This representation facilitated inverse design of inorganic crystals with targeted properties such as formation energy and bandgap, representing an important step toward property-conditioned generation.

Generative adversarial networks (GANs) [Goodfellow et al., 2014] have also been used for crystal structure prediction and candidate generation. In this framework, a discriminator and a generator neural network are trained, with the first being trained to distinguish between real data and generated samples, and the second trained adversarially to produce samples that fool the discriminator. Kim et al. [2020] used GANs to predict crystal structures conditioned on composition. It was notably one of the first methods to introduce a coordinate-based representation for crystal structure generation. Long et al. [2021] introduced a method to constrain generation to desired properties by additionally optimizing with respect to a property prediction network. CubicGAN [Zhao et al., 2021] introduced a learnable space group aware component to the representation, which proved influential, although restricted to cubic systems.

A.2 Diffusion and Flow-based Models

Diffusion models [Sohl-Dickstein et al., 2015, Ho et al., 2020, Song et al., 2021] have become a staple in image generation and have also been adapted successfully for crystal generation. They have shown promising capabilities in generating stable 3D periodic structures for novel crystalline materials. These approaches operate on the principle of gradually denoising randomly sampled configurations to produce physically plausible structures, incorporating both geometric and chemical constraints.

Early implementations such as CDVAE [Xie et al., 2021] and SyMat [Luo et al., 2023] combined VAEs with score-based denoising networks operating directly on atomic coordinates. These hybrid models first predict lattice parameters and atomic composition via the VAE component, and then predict atomic coordinates through score-based diffusion. Subsequent architectures like DiffCSP [Jiao et al., 2023] and MatterGen [Zeni et al., 2025] advanced this approach by adopting joint diffusion frameworks that simultaneously model atomic composition, fractional coordinates, and lattice parameters. These methods leverage SE(3)-equivariant neural networks adapted for periodicity [Duval et al., 2023] to predict the denoised composition, coordinates and lattice parameters, ensuring both Euclidean and periodic invariance in the learned distribution. An alternative representation strategy was proposed in UniMat [Yang et al., 2023], which introduced a unified crystal representation using a 4D tensor where atomic positions are stored within corresponding atom entries in the periodic table. This approach employs interleaved attention and convolution layers as denoising networks, effectively addressing the challenge of jointly modeling discrete atom types and continuous atomic coordinates.

Recent work has further extended diffusion models to incorporate domain knowledge and constraints. DiffCSP++ [Jiao et al., 2024] and SymmCD [Levy et al., 2025] enforce space group symmetry constraints during generation, significantly reducing the effective dimensionality of the generative task. By constraining atomic coordinates to follow specified Wyckoff positions and space group symmetries, these models not only

improve computational efficiency but also ensure that generated structures adhere to physically meaningful crystallographic rules. DiffCSP++ is based on using predefined structural templates for each space group, whereas in SymmCD the templates are also generated, resulting in increased diversity.

Text-guided diffusion models represent another frontier, with approaches like TGDMat [Das et al., 2025] and Chemeleon [Park et al., 2025] incorporating contextual representations from pretrained language models to enable generation of stable periodic materials that align with conditions specified in natural language descriptions.

Flow matching (FM) methods [Lipman et al., 2023, Albergo and Vanden-Eijnden, 2023, Liu et al., 2023] are closely related to diffusion models [Gao et al., 2025], and learn a time-dependent velocity field that smoothly transforms an arbitrary base distribution into the target distribution of stable materials. FlowMM [Miller et al., 2024] pioneered this approach for crystal generation, defining specialized representations that enforce global rotational and translational symmetries while respecting periodic boundary conditions. Flow matching notably offers faster generation than diffusion models, since it requires less evaluations of the neural network. Building on this work, FlowLLM [Sriram et al., 2024] introduced a hybrid approach that leverages language models to provide a chemically informed base distribution. Rather than starting from simple distributions, FlowLLM generates initial material representations using an LLM trained on known structures [Gruver et al., 2024], then refines these samples through flow matching that preserves composition while updating atomic positions and lattice parameters. Sampling from a base distribution defined by an LLM resulted in an increase in the rate of stable materials compared to sampling a generic base distribution.

Recent work demonstrates the potential of encoding both periodic (crystalline) and non-periodic (molecular) structures using a single unified flow matching framework. ADiT [Joshi et al., 2025] utilizes an autoencoder to learn a joint latent representation of both crystals and molecules, followed by flow matching in the unified latent space to generate new latent representations and decode them to valid crystals or molecules. This approach enables learning from and generating diverse materials, including metal-organic frameworks (MOFs), which exhibit both periodic and molecular characteristics. Unlike previous approaches that require specialized equivariant architectures, ADiT simplifies the generative process by using a standard Gaussian base distribution and Transformer architecture as the time-dependent velocity field in latent space. This simplification leads to faster generation times while still demonstrating strong performance on generation of stable structures.

A.3 Reinforcement Learning and GFlowNets

By contrast with generative models that learn from datasets of existing structures, reinforcement learning (RL) [Sutton and Barto, 1998, Arulkumaran et al., 2017] methods instead train an agent (i.e., a neural network) to optimize a given quantity or reward, without data. The model thus learns a behaviour policy by taking actions that act on a defined environment, like moving pieces (actions) on a chessboard (environment) to play a game, exploring actions that are likely to lead to high rewards (e.g., winning the game). This method has been used by Zamaraeva et al. [2023] with actions such as swapping atom positions and changing unit cell parameters, and the reward defined as the prediction of the negative change in energy from that action. Challenges however lie in access to accurate, but inexpensive energy calculation, as well as efficient exploration strategies. The reinforcement learning approach was extended to crystal design in work by Govindarajan et al. [2024]. In this work, actions consist of choosing new atoms to sequentially add to sites in a crystal. An offline RL framework is used, which allows to use a more expensive DFT-based method for computing rewards.

Generative Flow Networks (GFlowNets) [Bengio et al., 2021, 2023] offer a distinct set of methods, drawing inspiration from reinforcement learning [Tiapkin et al., 2024, Deleu et al., 2024]. GFlowNets formulate generation as a sequential decision process, constructing molecules or materials step-by-step using policy network(s) acting on pre-defined environments. In Crystal-GFN [Mila AI4Science et al., 2023], a space group is first sampled, then the atomic composition and finally the lattice parameters to form a crystal. It also encodes rigorously physico-chemical and symmetry constraints by masking impossible actions along the generation process. For materials discovery, the reward function typically corresponds to properties of interest such as formation energy, bandgap, or stress. Crystal-GFN demonstrated this capability by generating structures with low formation energy, targeted bandgaps, and high cell density, while other implementations have extended the approach to generate reticular frameworks for CO₂ capture [Cipcigan et al., 2024].

A.4 Large Language Models and Multimodal ML Systems

Large language models (LLMs) are increasingly being repurposed across the natural sciences as general-purpose priors for reasoning over sequences, graphs, and spatial data [Zhang et al., 2024]. This paradigm has recently been extended to materials generation. By representing crystal structures through textual descriptions of unit cells and atomic positions, language models trained to generate discrete tokens have shown competitive and sometimes superior performance compared to specialized geometric models [Gruver et al., 2024]. This unexpected effectiveness may stem from language models’ efficiency as information compressors [Delétang et al., 2023], enabling sample-efficient learning in data-constrained materials domains.

Early explorations by Flam-Shepherd and Aspuru-Guzik [2023] demonstrated that decoder-based transformers trained on text sequences could generate molecules, crystals, and protein binding sites. Building on this foundation, CrystalLLM [Antunes et al., 2024] fine-tuned language models on Crystallographic Information Files (CIF) for crystal structure generation. The approach was further refined by deCIFer [Johansen et al., 2025] and Gruver et al. [2024]; whereas deCIFer conditioned a decoder-only transformer on powder X-ray diffraction (PXRD) signals to generate crystal structures that agree with observed diffraction patterns in CIF format, Gruver et al. [2024] fine-tuned LLaMA-2 [Touvron et al., 2023] on condensed crystal format descriptions that included unit cell parameters and fractional atomic coordinates. This achieved great generation performance while enabling flexible natural language conditioning, supporting unconditional generation, property-targeted generation, and structure completion.

Matra-Genoa [De Breuck et al., 2025] is another recent decoder-based autoregressive transformer model designed for inorganic crystal generation. By decomposing crystal structures into an interpretable token-based vocabulary, it efficiently generates structures conditioned on target properties such as composition, space group, and formation energies. Further, Mishra et al. [2024] recently demonstrated that foundational materials domain LLMs (LLaMat) with additional training on semantic and syntactic tasks related to CIF (LLaMat-CIF) enables superior crystal generation with high diversity.

Domain knowledge integration has been explored in models like WyFormer [Kazeev et al., 2025b] and CrystFormer [Cao et al., 2024], which incorporate crystallographic priors of space groups and Wyckoff positions [Hahn et al., 1983, Zhu et al., 2024]. Further improvements have come through reinforcement learning techniques such as direct preference optimization (DPO) [Rafailov et al., 2023], demonstrated in PLaID [Xu et al., 2025a,b], which significantly improved stability and S.U.N. rates by learning from comparative evaluations of generated structures by machine learning interatomic potentials. Their work showcases the potential of leveraging reinforcement learning as a framework for optimizing future long-tail conditional generation tasks.

Multimodal approaches represent another frontier, attempting to integrate structural, spectroscopic, and textual information for more comprehensive materials modeling. Systems like MultiMat [Moro et al., 2025], MoMa [Wang et al., 2025], and LLM-Fusion [Boyar et al., 2025] demonstrate early efforts to combine representations across modalities, while agentic frameworks like A-Lab [Szymanski et al., 2023] point toward autonomous laboratories that can propose, execute, and refine materials synthesis guided by computational predictions.

B Evaluation Metrics for Materials Generation

Unconditional generation refers to the task of producing valid, stable crystal structures without targeting specific properties or constraints. The following metrics assess the fundamental quality of generated structures:

Fundamental Validity Metrics. These ensure the outputs are physically meaningful and chemically plausible. In different terms, they serve as a sanity check both for model development and inference time. Note that all metrics may not be relevant for every material system.

- **Charge Neutrality:** The total valence charge of all atoms must sum to zero:

$$\sum_{i=1}^N q_i = 0 \quad (1)$$

where q_i is the nominal oxidation state of atom i in the structure. For this to be calculated, the oxidation states of every atom in the structure must first be assigned. Here, we have developed a hierarchical structure for determining oxidation states and charge neutrality:

1. If all atoms are metals, each atom is assigned a nominal oxidation state of zero, and the structure is labeled as charge balanced. This metallic assignment is determined by using the Pymatgen "BVAnalyzer" class and the SMACT "metallicity-score" function with a cutoff of 0.7 [Davies et al., 2019]. If either of these are satisfied, the structure is labeled as a metal and therefore labeled charge-balanced.
 2. If all atoms are not metals, the Pymatgen "get-oxi-state-decorated-structure" function [Ong et al., 2013] is used to assign oxidation states and determine charge balance.
 3. However, the function used above can fail to find oxidation states for structures that are not well optimized. It is still necessary to determine whether these structures are charge-balanced, particularly in the case of generative model benchmarks, when many structures may be too far from typical structures for the Pymatgen functions to analyze them. Here, we determine charge neutrality using a data-driven approach from LeMat-Bulk [Siron et al., 2025]. First, this workflow determines all the possible charge-balanced compositions of oxidation states based on the observed oxidation states in LeMat-Bulk or the ICSD [Zagorac et al., 2019]. The most likely oxidation state assignments for this particular composition are determined by assigning each composition a score based on how probable that particular oxidation state configuration is, as determined by the distribution of oxidation states seen in LeMat-Bulk/ICSD. This score is determined by taking the geometric mean of all probabilities for each individual oxidation state. The geometric mean was chosen to normalize the score to the number of atoms in the structure, so scores between structures with two atoms and structures with ten atoms are comparable. Under the current implementation, if a composition exists from ICSD or LeMat-Bulk oxidation states that charge balances the structure, it passes the validity test. However, it is also possible to introduce a threshold here to filter out structures that require a combination of oxidation states which are highly uncommon, if a stricter validity metric is desired.
 4. If no charge-balanced composition can be made using oxidation states from LeMat-Bulk/ICSD, the structure is passed to the next stage of the validity assessment. In this stage, all historically observed oxidation states for each element, including highly irregular ones, are used to see if the structure can be charge-balanced. This list of oxidation states is drawn from Pymatgen's "Element" class. Since this list includes some very irregular oxidation states (negative charges on metals, for example), an additional check is used to determine whether this charge-balanced structure is viable. Since we don't have access to probabilities for these very uncommon states, the oxidation state of each element in the system is correlated to the Pauli electronegativity of that element. A negative correlation is desired, as this would show that increasing electronegativity leads to a decreasing oxidation state. Any structures that show a positive correlation are labeled invalid, as this would reflect a very chemically implausible structure overall.
- **Minimum Interatomic Distance:** All interatomic distances d_{ij} must exceed a cutoff value $d_{\min}$ to prevent atomic overlap. We use 0.7 Å in this work.

$$d_{ij} > d_{\min} \quad \forall i \neq j \quad (2)$$

- **Mass density and atomic number density** should be within reasonable ranges. Mass density is given by $\rho = \frac{M_{\text{total}}}{V_{\text{cell}}}$, in (g/cm^3). The latter is expressed in atoms/Å³. We take upper bounds of 25 g/cm³ and 0.5 atoms/Å³, respectively.
- **Valid crystallographic representation** is a good proxy to determine whether a structure is CIF-readable using *pymatgen*.
- **Lattice parameters** should be within reasonable ranges. We take upper bounds of 100 Å for a,b,c, and 180 degrees for α , β , γ respectively, and lower bounds of 1 Å and zero degrees for a,b,c and α , β , γ , respectively.




Figure 7 The elemental composition of the structures labeled invalid by SMACT yet labeled valid by our model across the 5 random samples of 1000 structures drawn from LeMat-Bulk. Of these 5000 structures, 631 are labeled valid by our model but invalid by SMACT. The counts indicate the number of times a structure containing that element is found in the overall list of structures. For example, 70 of the 631 structures contain boron.




Figure 8 Seed variation in validity outcomes. Five independent samples of 1,000 structures each from LeMat-Bulk. Categories match the agreement breakdown from Figure 3. Dashed lines indicate mean values. High consistency across seeds demonstrates robustness of both validation methods, with our approach consistently validating ~ 126 additional structures per seed while rejecting only ~ 10 SMACT-approved structures.

Stability metrics. These assess the thermodynamic and energetic properties of generated structures:

- **Formation Energy (E_f):**

$$E_f = E_{\text{tot}}(\text{compound}) - \sum_i n_i \mu_i \quad (3)$$

where E_{tot} is the total energy of the crystal, n_i is the number of atoms of element i , and μ_i is the chemical potential of the pure element. The result is normalized per atom: $E_f^{\text{per atom}} = \frac{E_f}{\sum_i n_i}$. We want it to be as small (and negative) as possible.

The chemical potentials μ_i are derived from the LeMaterial-Bulk dataset using the same methodology as used in Ong et al. [2008].

Multi-MLIP Ensemble Implementation: The formation energy metric supports ensemble statistics across multiple MLIPs (ORB, MACE, UMA). For each structure, ensemble statistics are computed as:

$$\langle E_f \rangle = \frac{1}{N_{\text{MLIP}}} \sum_{k=1}^{N_{\text{MLIP}}} E_f^{(k)} \quad (4)$$

$$\sigma_{E_f} = \sqrt{\frac{1}{N_{\text{MLIP}} - 1} \sum_{k=1}^{N_{\text{MLIP}}} \left(E_f^{(k)} - \langle E_f \rangle \right)^2} \quad (5)$$

where $E_f^{(k)}$ is the formation energy predicted by the k -th MLIP. The implementation extracts pre-computed ensemble statistics from structure properties (`formation_energy_mean`, `formation_energy_std`) or calculates them from individual MLIP results (`formation_energy_orb`, `formation_energy_mace`, etc.). A minimum of 2 MLIPs is required for ensemble statistics.

- **Energy Above Convex Hull (E_{hull}):**

$$E_{\text{hull}} = E_{\text{tot}} - E_{\text{hull}}^{\min} \quad (6)$$

Structures with $E_{\text{hull}} \leq 0$ are considered stable, while values below approximately 0.1 eV/atom are often deemed metastable. We take LeMat-Bulk [Siron et al., 2024] as reference point for calculating the convex hull.

The convex hull is constructed by filtering the LeMat-Bulk dataset to include only compounds containing elements present in the target composition, creating `PEntry` objects, and using `Pymatgen's PhaseDiagram.get_decomp_and_e_above_hull()` method. The implementation handles charged species by extracting neutral elements before phase diagram construction. Multi-MLIP ensemble statistics follow the same formulation as formation energy: $\langle E_{\text{hull}} \rangle = \frac{1}{N_{\text{MLIP}}} \sum_{k=1}^{N_{\text{MLIP}}} E_{\text{hull}}^{(k)}$ with corresponding standard deviation calculations.

- **Relaxation Stability:** We use an ensemble of Machine Learning Interatomic Potentials to relax the generated structures (each one is done independently). Then, compute the Root Mean Square Deviation (RMSD) between pre- and post-relaxation atomic positions:

$$\text{RMSD} = \sqrt{\frac{1}{N} \sum_{i=1}^N \|\mathbf{r}_i^{\text{init}} - \mathbf{r}_i^{\text{relax}}\|^2} \quad (7)$$

Low RMSD indicates minimal distortion and structural robustness under optimization.

The implementation calculates individual RMSD values for each MLIP relaxation, then computes ensemble statistics: $\langle \text{RMSD} \rangle = \frac{1}{N_{\text{MLIP}}} \sum_{k=1}^{N_{\text{MLIP}}} \text{RMSD}^{(k)}$ where $\text{RMSD}^{(k)}$ is the relaxation RMSD from the k -th MLIP. The metric extracts pre-computed values from structure properties (`relaxation_rmsd_mean`, `relaxation_rmsd_std`) or calculates ensemble statistics from individual MLIP results (`relaxation_rmsd_orb`, `relaxation_rmsd_mace`, etc.). We use a force convergence threshold of 0.02 eV/Å and a maximum of 50 optimization steps for reporting RMSD in this work.

Novelty, Uniqueness, and Diversity Metrics. These evaluate how effectively a model explores the chemical space:

- **Novelty ($\mathcal{N}$):** Evaluates the fraction of generated structures that are not present in a reference dataset of known materials. The novelty score is defined as:

$$\mathcal{N} = \frac{|\{x \in G \mid x \notin T\}|}{|G|} \quad (8)$$

where G is the set of generated structures and T is the reference dataset (LeMat-Bulk).

The implementation supports two comparison methods: **BAWL fingerprinting** using crystallographic hash strings with the Weisfeiler-Lehman algorithm, and **structure matching** using Pymatgen’s symmetry-aware structural comparison algorithms. For BAWL, novelty is determined by checking if the generated structure’s fingerprint exists in the pre-computed reference fingerprint set. For structure matching, each generated structure is compared against reference structures with overlapping elemental compositions using space group analysis and atomic position matching with configurable tolerances. In our paper, we report results using *pymatgen’s StructureMatcher* approach for more robust structural comparison against the LeMat-Bulk reference dataset.

- **Uniqueness ($\mathcal{U}$):** Measures the fraction of unique structures within the generated set to assess internal diversity. The uniqueness score is defined as:

$$\mathcal{U} = \frac{|\text{unique}(G)|}{|G|} \quad (9)$$

where $\text{unique}(G)$ returns the set of unique structures based on their fingerprints.

The metric is implemented as a structure-level continuous scoring system rather than binary classification. For BAWL fingerprinting, individual uniqueness scores are assigned as $u_i = 1/c_i$, where c_i is the count of structures sharing the same fingerprint within the generated set. This assigns a score of 1.0 to truly unique structures while proportionally penalizing duplicated structures. For structure matching, the implementation uses pairwise comparison with an ordered approach: structure i is considered unique if it is not equivalent to any structure j where $j < i$, ensuring deterministic selection of the first occurrence as the unique representative. The overall uniqueness metric is computed as $\mathcal{U} = \frac{1}{|G|} \sum_{i=1}^{|G|} u_i$. Both BAWL fingerprinting and structure matching methods are supported, with structure matching used for reporting results.

- **S.U.N. and M.S.U.N. Rates:** Proportion of generated structures that are simultaneously Stable (or Metastable), Unique, and Novel:

$$\text{S.U.N. Rate} = \frac{|\text{novel}(\text{unique}(x \in G \mid E_{\text{hull}}(x) \leq 0))|}{|G|} \quad (10)$$

$$\text{M.S.U.N. Rate} = \frac{|\text{novel}(\text{unique}(x \in G \mid 0 < E_{\text{hull}}(x) \leq \tau))|}{|G|} \quad (11)$$

The implementation follows a hierarchical computation order: **Stability** $\rightarrow$ **Uniqueness** $\rightarrow$ **Novelty**. First, structures are classified as stable ($E_{\text{hull}} \leq 0$) or metastable ($0 < E_{\text{hull}} \leq \tau$) using energy above hull values computed by the Multi-MLIP stability preprocessor. Then, uniqueness is evaluated within each stability class separately using the chosen comparison method. Finally, novelty is assessed for unique structures from each stability class. This hierarchical approach provides detailed metrics at each evaluation stage: stability counts, unique-within-stable/metastable counts, and final SUN/MSUN counts. The Multi-MLIP preprocessor assigns ensemble stability properties (e.g., `e_above_hull_mean`) to structure objects, enabling robust stability classification across multiple MLIPs (ORB, MACE, UMA). We set τ to **0.1 eV/atom** for assembling results.

- **Diversity:** We assess diversity by comparing the distribution of generated structures against reference datasets across key crystallographic and compositional attributes. This includes space groups, elemental compositions, lattice parameters, and atomic environments. In addition to distribution plots, we compute feature-level entropy metrics to quantify the spread of generated samples.
 - **Composition, Space Group, Lattice, and Atomic Site Entropy:** Given a set of N generated structures, we first compute the frequency f_i of each category (e.g., element i for composition,

space group index for symmetry, discretized lattice bins, or Wyckoff site labels). These frequencies are normalized into a probability distribution $p_i = \frac{f_i}{\sum_j f_j}$. We then compute the Shannon entropy

$$H = - \sum_i p_i \log p_i,$$

together with the Vendi score [Friedman and Dieng, 2022], defined as the exponential of the Shannon entropy. Although illustrated here for composition entropy, the same procedure is applied to the other diversity attributes in our benchmark.

Distribution-Level Metrics. To evaluate how well the distribution of generated structures matches that of real materials, we use the following metrics:

- **Jensen–Shannon Distance** [Fuglede and Topsoe, 2004]:

$$\text{JSD}(P, Q) = \sqrt{\frac{1}{2}D_{KL}(P \parallel M) + \frac{1}{2}D_{KL}(Q \parallel M)} \quad (12)$$

where P and Q are the distributions of generated and real samples, respectively, $M = \frac{1}{2}(P + Q)$ is their mixture distribution, and D_{KL} denotes the Kullback–Leibler divergence.

- **Maximum Mean Discrepancy (MMD)** [Tolstikhin et al., 2016]:

$$\text{MMD}^2(P, Q) = \mathbb{E}_{x, x'}[k(x, x')] + \mathbb{E}_{y, y'}[k(y, y')] - 2 \mathbb{E}_{x, y}[k(x, y)] \quad (13)$$

where P and Q are distributions of generated and real samples, and k is a kernel function.

- **Fréchet Distance Metrics** [Heusel et al., 2017, Preuer et al., 2018]: Variants such as the Fréchet ChemNet Distance (FCD) compare the distributions of generated and reference structures:

$$\text{FD}(G, T) = \|\mu_G - \mu_T\|^2 + \text{Tr}\left(\Sigma_G + \Sigma_T - 2(\Sigma_G \Sigma_T)^{1/2}\right) \quad (14)$$

where μ and Σ denote the mean and covariance in embedding space.

Model Efficiency. This measures how effectively a model learns from limited training data [Gao et al., 2022]:

- **Generic Metrics:** training dataset size, number of model parameters, number of epochs required for training, training time and associated computational infrastructure, inference time on 10k structures.
- **Learning Curve Analysis:** Performance (e.g., S.U.N. rate, property prediction accuracy) as a function of the number of expensive function evaluations (e.g., DFT calculations) required for training, i.e., the number of labeled data points.

Herfindahl-Hirschman Index (HHI) Metrics. The Herfindahl-Hirschman index quantifies supply risk concentration for materials by measuring the concentration of element production sources and reserves. The elemental data for calculating HHI metrics are derived from the calculator⁸ made available by Mansouri Tehrani et al. [2017]. For a given crystal structure with composition, we compute:

- **Compound HHI Value:** For a compound with chemical formula represented by composition C :

$$\text{HHI}_{\text{compound}} = \sum_i x_i \cdot \text{HHI}_i \quad (15)$$

where x_i is the fractional composition of element i in the compound, and HHI_i is the element-specific HHI value.

⁸<https://sparks.mse.utah.edu/about.html>

- **Production HHI:** Measures supply risk based on the concentration of element production sources (market concentration):

$$\text{HHI}_{\text{production}} = \sum_j s_j^2 \times 10000 \quad (16)$$

where s_j is the market share of producer j for a given element.

- **Reserve HHI:** Measures long-term supply risk based on the concentration of element reserves (geographic distribution):

$$\text{HHI}_{\text{reserve}} = \sum_k r_k^2 \times 10000 \quad (17)$$

where r_k is the fraction of global reserves held by country/region k .

- **Scaling Convention:** HHI values are typically scaled from the classical range $[0, 10000]$ to a convenience range $[0, 10]$:

$$\text{HHI}_{\text{scaled}} = \frac{\text{HHI}_{\text{classical}}}{1000} \quad (18)$$

- **Combined HHI Score:** The final benchmark score combines both production and reserve metrics using weighted averaging:

$$\text{HHI}_{\text{combined}} = w_{\text{prod}} \cdot \text{HHI}_{\text{production}} + w_{\text{res}} \cdot \text{HHI}_{\text{reserve}} \quad (19)$$

where $w_{\text{prod}} = 0.25$ and $w_{\text{res}} = 0.75$ by default, prioritizing long-term supply security over short-term market dynamics.

- **Missing Element Handling:** Elements not found in the HHI lookup tables are assigned the maximum risk value (10000 unscaled / 10 scaled) to represent maximum supply uncertainty for rare or untracked elements.
- **Risk Categories:** For the scaled $[0, 10]$ range:

$$\text{Low Risk : } \text{HHI}_{\text{scaled}} \leq 2.0 \quad (20)$$

$$\text{Moderate Risk : } 2.0 < \text{HHI}_{\text{scaled}} \leq 5.0 \quad (21)$$

$$\text{High Risk : } \text{HHI}_{\text{scaled}} > 5.0 \quad (22)$$

C Benchmark Results

Table 3 Training datasets and data sources used for the reported generative crystal structure models.

Model	Training Dataset	Source of Submitted Structures
ADiT [Joshi et al., 2025]	MP-20	Authors of [Joshi et al., 2025]
Crystal GFN [AI4Science et al., 2023]	MP-20	Authors of [AI4Science et al., 2023]
Crystalformer[Cao et al., 2024]	MP-20	Figshare of [Kazeev et al., 2025a] ⁹
DiffCSP [Jiao et al., 2023]	MP-20	Figshare of [Kazeev et al., 2025a] ⁹
DiffCSP++ [Jiao et al., 2024]	MP-20	Figshare of [Kazeev et al., 2025a] ⁹
LLaMat2 [Mishra et al., 2024]	MP-20	Authors of [Mishra et al., 2024]
MatterGen [Zeni et al., 2025]	MP-20	Figshare of [Kazeev et al., 2025a] ⁹
PLaID++ [Xu et al., 2025a]	MP-20	Authors of [Xu et al., 2025a]
SymmCD [Levy et al., 2025]	MP-20	Figshare of [Kazeev et al., 2025a] ⁹
WyFormer-DiffCSP++ [Kazeev et al., 2025a]	MP-20	Authors of [Kazeev et al., 2025a]
WyFormer-DiffCSP++-DFT [Kazeev et al., 2025a]	MP-20	Authors of [Kazeev et al., 2025a]
Alexandria [Schmidt et al., 2021, 2024]	—	Sampled from source database
OQMD [Saal et al., 2013, Kirklin et al., 2015]	—	Sampled from source database
AFLOW [Curtarolo et al., 2012]	—	Sampled from source database

As summarized in Table 3, since most models included in this study were trained on MP-20, we report results using the multi-MLIP ensemble benchmark on both the LeMat-Bulk and MP-20 datasets as references. Evaluating on both reference sets ensures consistency between the training distribution and the evaluation

⁹https://figshare.com/articles/dataset/Generated_crystals_for_WyFormer_DiffCSP_DiffCSP_WyCryst_SymmCD_CrystalFormer_MiAD/29145101

database, allowing authors to better interpret model behavior without being penalized for discrepancies arising solely from dataset choice.

However, LeMat-Bulk is currently the most comprehensive and diverse collection of inorganic crystal structures. Performance on this dataset, therefore, provides the most informative assessment of generative modeling capability and serves as the primary reference for the public leaderboard. Metrics that depend on a reference dataset are expected to vary between the two settings, including novelty scores, energy-based metrics, stability and metastability rates, and embedding-based distributional metrics such as the Fréchet distance.

Several generative models considered here perform structure relaxation as part of their generation pipeline (MatterGen, PLaID++, WyFormer, and WyFormer-DFT). To provide a fair and transparent comparison, we report two sets of results for all models and baselines: (i) metrics computed directly on the submitted structures, and (ii) metrics computed after relaxing the structures using the UMA s-1 potential. In this setting, only UMA s-1 is used for evaluating formation energies and constructing the convex hull. Using the full multi-MLIP ensemble would not be meaningful, since one of the potentials is already used for relaxation and the goal is to measure geometric and energetic quality in a self-consistent manner.

For all relaxation procedures, we use the default configuration of our relaxation pipeline: a force convergence threshold of 0.02 eV/Å and a maximum of 500 optimization steps. These settings correspond to the default parameters of the relaxation module:

```
--fmax 0.02
--steps 500
```

This ensures that all post-relaxation metrics are evaluated under a uniform and reproducible protocol across models and datasets.

C.1 Model Training Details

Table 4 Training and inference times for 2,500 samples for the reported generative models, with respective compute information.

Model	Training Time (hrs)	Train GPU	Inference Time (s)	Inference GPU
SymmCD	26	V100	1,400	RTX8000
Crystalformer	13	A100	100	V100
ADiT	576	V100	300	V100
PLaID++	10	H100	345	H100
MatterGen	80	A100	18,000	V100
DiffCSP	–	–	2,335	A40
DiffCSP++	19.5	RTX 6000	6,150	A40
WyFormer	11.5	RTX 6000	0.625	RTX6000
Crystal-GFN	12	CPU	3000	CPU
LLaMat2	3264	A100	3897	A100
LLaMat3	144	CS2 Cerebras	4235	A100


C.2 Benchmark Results with LeMat-Bulk as the reference database

We visualize S.U.N. and M.S.U.N. structures with respect to LeMat-Bulk from various generative models using t-SNE, a dimensionality reduction method in Figure 9. The resulting plots show that most models produce samples that are broadly dispersed throughout crystallographic space, rather than collapsing into small or isolated clusters. This indicates that many contemporary crystal generators are capable of proposing chemically diverse structures. One standout model in these plots is MatterGen, whose samples appear to have the highest coverage of ORB’s embedding space.


We report the model efficiency metrics in Figure 10. Our pareto frontier demonstrates that there does exist a correlation between model parameters and general M.S.U.N. performance, however, careful architectural design is still critical. On the inference-time axis, WyFormer occupies the extreme low-latency regime while MatterGen sits at the opposite end, delivering the highest M.S.U.N.% at the expense of the longest runtimes. While the reported runtimes are not perfectly comparable—since models were evaluated on heterogeneous hardware—they nonetheless highlight the importance of prioritizing architectures that deliver

strong performance under realistic computational budgets, which we view as critical for scalable downstream discovery.

A similar picture emerges when comparing M.S.U.N.% against the number of model parameters. WyFormer and Crystalformer achieve competitive M.S.U.N.% with comparatively small parameter counts, and DiffCSP and MatterGen extend this frontier to larger models. In contrast, SymmCD, ADiT, and especially LLaMat2/3 use one to two orders of magnitude more parameters than the frontier models yet deliver comparable or worse M.S.U.N.%. Together, these results suggest that architectures explicitly tailored to crystal generation (WyFormer, Crystalformer, DiffCSP, and MatterGen) provide better M.S.U.N.%-per-parameter and M.S.U.N.%-per-second than more generic, heavily overparameterized generators.




(a) t-SNE visualization of ORB graph embeddings for M.S.U.N. materials with LeMatBulk as the reference set.




(b) t-SNE visualization of ORB graph embeddings for S.U.N. materials with LeMatBulk as the reference set.

Figure 9 Two-dimensional t-SNE visualizations (perplexity = 30) of ORB graph embeddings of generative model structures computed with LeMatBulk as the reference set. The visualizations demonstrate how MatterGen, DiffCSP, and PlaID++ generate a wide diversity of novel crystal structures which cover a breadth of chemical space.



(a) Inference time vs MSUN%



(b) Model parameters vs MSUN%

Figure 10 Performance-efficiency trade-offs for crystal structure generation models. (a) Inference time (log scale) versus MSUN% shows the computational cost required to achieve MSUN structures. Models on the Pareto frontier (connected by blue line) represent optimal time-performance trade-offs. (b) Model parameter count (log scale) versus MSUN% illustrates the relationship between model size and generation quality. WyFormer, Crystalformer, DiffCSP, and MatterGen form the Pareto frontier, demonstrating that larger models can achieve better performance when properly designed.

Table 5 Model evaluation metrics calculated using **MLIP ensemble** with **LeMat-Bulk** as the reference set: Validity, Energy, Stability, and Metastability. All models provide 2500 structures except Crystalformer (1000 structures). Submissions are organized into four categories separated by horizontal lines: (1) models that incorporate relaxation as part of their generation pipeline (MatterGen, PLaID++, WyFormer, WyFormer-DFT), (2) other generative models, (3) samples from real-world datasets that contribute to LeMat-Bulk (Alexandria, OQMD) as baselines, and (4) samples from the AFLOW dataset and MP-Exp structures. Best values are shown in **bold**, second-best values are underlined. Arrows indicate whether higher ($\uparrow$) or lower ($\downarrow$) values are better.

Model	Validity				Unique% $\uparrow$	Novel% $\uparrow$	Energy-based			Stability			Metastability		
	Valid% $\uparrow$	CN% $\uparrow$	MinDist% $\uparrow$	PhysPlau% $\uparrow$			FormE (Std) $\downarrow$	E_{hull} (Std) $\downarrow$	RMSD (Std) $\downarrow$	Stable% $\uparrow$	U-Stable% $\uparrow$	SUN% $\uparrow$	Metastable% $\uparrow$	U-Meta% $\uparrow$	MSUN% $\uparrow$
MatterGen	95.7	95.9	99.8	100.0	95.1	70.5	-0.70 ± 0.79	0.18 ± 0.18	0.39 ± 0.50	2.0 (49)	2.0 (49)	0.2 (6)	33.4 (835)	32.9 (823)	15.0 (374)
PLaID++	96.0	96.0	96.2	96.2	77.8	24.2	-0.50 ± 0.44	0.09 ± 0.16	0.13 ± 0.29	12.4 (311)	8.6 (215)	1.0 (26)	60.7 (1517)	47.2 (1181)	7.6 (189)
WyFormer	93.4	95.2	98.2	100.0	93.0	<u>66.4</u>	-0.43 ± 0.51	0.50 ± 0.51	0.81 ± 0.38	0.5 (12)	0.1 (3)	0.1 (3)	15.7 (393)	15.5 (387)	1.9 (48)
WyFormer-DFT	95.2	95.2	100.0	100.0	<u>95.0</u>	<u>66.4</u>	-0.67 ± 0.91	0.27 ± 0.36	0.42 ± 0.60	<u>3.7</u> (92)	<u>3.5</u> (88)	<u>0.4</u> (9)	24.8 (621)	24.8 (619)	7.8 (195)
ADiT	90.6	99.0	91.5	100.0	87.8	26.0	-0.73 ± 0.93	0.33 ± 0.45	0.38 ± 0.40	0.4 (10)	0.4 (10)	0.0 (0)	<u>36.5</u> (913)	<u>35.0</u> (874)	1.0 (25)
Crystal-GFN	51.7	58.3	91.1	96.3	51.7	51.7	-1.30 ± 2.63	2.09 ± 2.38	1.87 ± 0.97	0.0 (0)	0.0 (0)	0.0 (0)	0.0 (0)	0.0 (0)	0.0 (0)
Crystalformer	69.9	96.3	72.9	99.9	69.4	31.8	-0.17 ± 1.48	0.70 ± 1.33	0.66 ± 1.05	1.4 (14)	1.4 (14)	0.0 (0)	28.8 (288)	28.5 (285)	3.1 (31)
DiffCSP	<u>95.7</u>	96.2	99.5	100.0	94.8	66.2	-0.64 ± 0.96	0.28 ± 0.56	0.59 ± 0.63	2.3 (58)	2.2 (55)	0.1 (3)	29.8 (745)	29.0 (726)	<u>8.5</u> (212)
DiffCSP++	95.3	95.5	99.7	100.0	95.1	62.0	-0.52 ± 0.87	0.41 ± 0.44	0.69 ± 0.82	1.0 (25)	1.0 (25)	0.2 (4)	26.4 (660)	26.2 (655)	5.0 (124)
LLaMat2	84.4	<u>97.0</u>	87.1	99.6	81.4	30.0	-0.47 ± 1.01	0.44 ± 0.71	0.54 ± 0.71	0.7 (18)	0.7 (18)	0.1 (3)	34.7 (867)	32.4 (811)	2.1 (52)
LLaMat3	15.4	82.5	16.8	84.2	15.2	10.5	0.74 ± 2.69	1.71 ± 2.48	1.01 ± 1.10	0.1 (2)	0.1 (2)	0.0 (0)	2.1 (53)	2.0 (51)	0.2 (4)
SymmCD	73.4	95.9	76.4	100.0	73.0	47.0	-0.02 ± 1.22	0.88 ± 1.07	0.87 ± 1.09	1.4 (34)	1.4 (34)	0.1 (2)	18.6 (464)	18.3 (457)	2.4 (59)
Alexandria	93.4	93.4	99.0	99.0	93.4	0.0	-0.18 ± 0.94	0.44 ± 0.61	0.14 ± 0.32	2.3 (57)	2.3 (57)	0.0 (0)	27.4 (684)	27.4 (684)	0.0 (0)
OQMD	96.8	96.8	99.4	99.5	96.4	0.0	-0.23 ± 0.88	0.39 ± 0.46	0.14 ± 0.28	5.3 (132)	5.3 (132)	0.0 (0)	29.1 (727)	29.0 (724)	0.0 (0)
AFLOW	91.4	91.4	100.0	100.0	91.4	21.5	-0.19 ± 0.86	0.35 ± 0.42	0.12 ± 0.44	9.3 (232)	8.6 (216)	0.0 (0)	30.4 (761)	27.6 (690)	0.7 (17)
MP-Exp	98.0	98.0	100.0	100.0	85.6	1.9	-1.02 ± 0.97	0.09 ± 0.45	0.20 ± 0.35	15.7 (393)	13.6 (339)	0.3 (8)	64.1 (1602)	56.5 (1412)	0.6 (14)

Table 6 Model evaluation metrics calculated using **MLIP ensemble** with **LeMat-Bulk** as the reference set: Distribution, Diversity, and HHI. All models provide 2500 structures except Crystalformer (1000 structures). Submissions are organized into four categories separated by horizontal lines: (1) models that incorporate relaxation as part of their generation pipeline (MatterGen, PLaID++, WyFormer, WyFormer-DFT), (2) other generative models, (3) samples from real-world datasets that contribute to LeMat-Bulk (Alexandria, OQMD) as baselines, and (4) samples from the AFLOW dataset and MP-Exp structures. Best values are shown in **bold**, second-best values are underlined. Arrows indicate whether higher ($\uparrow$) or lower ($\downarrow$) values are better.

Model	Distribution			Diversity				HHI	
	JS $\downarrow$	MMD $\downarrow$	FID $\downarrow$	ElemDiv $\uparrow$	SGDiv $\uparrow$	SizeDiv $\uparrow$	SiteDiv $\uparrow$	Prod $\downarrow$	Res $\downarrow$
MatterGen	0.44	0.006	2.13	0.64	0.14	0.24	12.78	3.52	2.72
PLaID++	0.43	0.029	2.69	0.68	0.26	0.21	6.01	5.23	3.36
WyFormer	<u>0.22</u>	0.006	1.55	<u>0.74</u>	<u>0.50</u>	0.27	<u>23.84</u>	3.55	2.73
WyFormer-DFT	0.25	0.010	1.72	0.73	0.49	0.26	23.56	3.49	2.74
ADiT	0.51	0.002	1.72	0.71	0.03	0.23	14.18	3.52	2.78
Crystal-GFN	0.55	0.121	43.17	0.75	0.29	0.32	32.06	<u>3.48</u>	2.34
Crystalformer	0.26	<u>0.003</u>	2.31	0.73	0.35	<u>0.31</u>	18.59	3.78	2.82
DiffCSP	0.45	0.008	2.25	0.71	0.14	0.24	14.42	3.29	2.60
DiffCSP++	<u>0.21</u>	<u>0.003</u>	<u>1.28</u>	0.72	0.51	0.27	20.84	3.56	2.77
LLaMat2	0.32	0.002	1.23	0.71	0.25	0.24	9.86	3.94	2.86
LLaMat3	0.34	0.005	9.48	0.70	0.10	0.29	13.15	3.83	2.81
SymmCD	<u>0.22</u>	0.006	1.65	0.73	0.49	0.27	19.45	3.54	2.74
Alexandria	0.09	0.000	1.31	0.71	0.50	0.25	13.50	4.02	3.33
OQMD	0.27	0.004	0.83	0.63	0.48	0.24	12.57	4.22	3.29
AFLOW	0.28	0.004	1.78	0.69	0.46	0.25	10.07	4.12	3.22
MP-Exp	0.33	0.021	2.90	0.72	0.71	0.28	54.79	2.78	2.27

Table 7 Model evaluation metrics calculated using **uma** with **LeMat-Bulk** as the reference set: Validity, Energy, Stability, and Metastability. All models provide 2500 structures except Crystalformer (1000 structures). Submissions are organized into four categories separated by horizontal lines: (1) models that incorporate relaxation as part of their generation pipeline (MatterGen, PLaID++, WyFormer, WyFormer-DFT), (2) other generative models, (3) samples from real-world datasets that contribute to LeMat-Bulk (Alexandria, OQMD) as baselines, and (4) samples from the AFLOW dataset and MP-Exp structures. Best values are shown in **bold**, second-best values are underlined. Arrows indicate whether higher ($\uparrow$) or lower ($\downarrow$) values are better.

Model	Validity				Unique% $\uparrow$	Novel% $\uparrow$	Energy-based			Stability			Metastability		
	Valid% $\uparrow$	CN% $\uparrow$	MinDist% $\uparrow$	PhysPlau% $\uparrow$			FormE (Std) $\downarrow$	E_{hull} (Std) $\downarrow$	RMSD (Std) $\downarrow$	Stable% $\uparrow$	U-Stable% $\uparrow$	SUN% $\uparrow$	Metastable% $\uparrow$	U-Meta% $\uparrow$	MSUN% $\uparrow$
MatterGen	95.7	95.9	99.8	100.0	95.1	70.5	-0.69 ± 0.80	0.19 ± 0.19	0.26 ± 0.30	1.3 (32)	1.2 (31)	0.1 (3)	32.7 (818)	32.2 (806)	14.2 (354)
PLaID++	96.0	96.0	96.2	96.2	77.8	24.2	-0.45 ± 0.48	0.09 ± 0.17	0.10 ± 0.20	10.0 (250)	6.9 (173)	1.0 (24)	62.1 (1552)	47.8 (1194)	7.4 (185)
WyFormer	93.4	95.2	98.2	100.0	93.0	<u>66.4</u>	-0.41 ± 0.97	0.51 ± 0.52	0.53 ± 0.54	0.5 (12)	0.5 (12)	0.0 (1)	15.2 (381)	15.0 (375)	2.1 (53)
WyFormer-DFT	95.2	95.2	100.0	100.0	<u>95.0</u>	<u>66.4</u>	-0.65 ± 1.07	0.28 ± 0.65	0.19 ± 0.28	2.2 (56)	2.2 (55)	0.2 (5)	25.9 (648)	25.7 (643)	7.8 (195)
ADiT	90.6	99.0	91.5	100.0	87.8	26.0	-0.72 ± 0.95	0.34 ± 0.45	0.31 ± 0.39	0.1 (3)	0.1 (3)	0.0 (0)	35.4 (886)	33.9 (848)	0.9 (23)
Crystal-GFN	51.7	58.3	91.1	96.3	51.7	51.7	-0.56 ± 2.46	2.83 ± 2.23	0.83 ± 0.34	0.0 (0)	0.0 (0)	0.0 (0)	0.0 (0)	0.0 (0)	0.0 (0)
Crystalformer	69.9	96.3	72.9	99.9	69.4	31.8	-0.16 ± 1.48	0.71 ± 1.31	0.46 ± 0.61	1.4 (14)	1.4 (14)	0.0 (0)	28.4 (284)	28.1 (281)	3.2 (32)
DiffCSP	<u>95.7</u>	96.2	99.5	100.0	94.8	66.2	-0.62 ± 0.98	0.28 ± 0.58	0.37 ± 0.38	2.0 (49)	1.8 (46)	0.2 (4)	29.7 (743)	29.1 (727)	8.2 (205)
DiffCSP++	95.3	95.5	99.7	100.0	95.1	62.0	-0.50 ± 0.88	0.42 ± 0.44	0.47 ± 0.50	1.2 (30)	1.2 (30)	0.2 (4)	25.4 (634)	25.2 (629)	4.8 (120)
LLaMat2	84.4	<u>97.0</u>	87.1	99.6	81.4	30.0	-0.45 ± 1.03	0.45 ± 0.72	0.43 ± 0.49	1.0 (26)	1.0 (25)	0.2 (4)	33.4 (834)	31.2 (781)	1.9 (48)
LLaMat3	15.4	82.5	16.8	84.2	15.2	10.5	0.79 ± 2.73	1.74 ± 2.52	0.75 ± 0.84	0.0 (0)	0.0 (0)	0.0 (0)	2.2 (54)	2.1 (52)	0.2 (5)
SymmCD	73.4	95.9	76.4	100.0	73.0	47.0	0.00 ± 1.23	0.89 ± 1.08	0.58 ± 0.61	1.2 (29)	1.2 (29)	0.1 (2)	18.6 (460)	18.4 (460)	2.4 (60)
Alexandria	94.3	94.3	100.0	100.0	94.3	0.0	-0.16 ± 0.97	0.45 ± 0.64	0.11 ± 0.39	1.2 (29)	1.2 (29)	0.0 (0)	28.1 (695)	28.1 (695)	0.0 (0)
OQMD	97.3	97.3	100.0	100.0	96.9	0.0	-0.21 ± 0.89	0.40 ± 0.46	0.13 ± 0.30	3.6 (90)	3.6 (90)	0.0 (0)	30.6 (762)	30.5 (759)	0.0 (0)
AFLOW	91.4	91.4	100.0	100.0	87.1	21.5	-0.18 ± 0.86	0.35 ± 0.42	0.11 ± 0.40	5.8 (146)	5.4 (136)	0.0 (0)	33.5 (838)	30.4 (760)	0.7 (17)
MP-Exp	98.0	98.0	100.0	100.0	85.6	1.9	-0.99 ± 1.36	0.11 ± 1.05	0.19 ± 0.39	8.8 (221)	7.4 (186)	0.4 (9)	70.1 (1753)	61.8 (1545)	0.5 (13)

Table 8 Model evaluation metrics calculated using **uma** with **LeMat-Bulk** as the reference set: Distribution, Diversity, and HHI. All models provide 2500 structures except Crystalformer (1000 structures). Submissions are organized into four categories separated by horizontal lines: (1) models that incorporate relaxation as part of their generation pipeline (MatterGen, PLaID++, WyFormer, WyFormer-DFT), (2) other generative models, (3) samples from real-world datasets that contribute to LeMat-Bulk (Alexandria, OQMD) as baselines, and (4) samples from the AFLOW dataset and MP-Exp structures. Best values are shown in **bold**, second-best values are underlined. Arrows indicate whether higher ($\uparrow$) or lower ($\downarrow$) values are better.

Model	Distribution			Diversity				HHI	
	JS $\downarrow$	MMD $\downarrow$	FID $\downarrow$	ElemDiv $\uparrow$	SGDiv $\uparrow$	SizeDiv $\uparrow$	SiteDiv $\uparrow$	Prod $\downarrow$	Res $\downarrow$
MatterGen	0.44	0.006	0.38	0.64	0.14	0.24	12.78	3.52	2.72
PLaID++	0.43	0.029	0.96	0.68	0.26	0.21	6.01	5.23	3.36
WyFormer	<u>0.22</u>	0.006	<u>0.10</u>	<u>0.74</u>	<u>0.50</u>	<u>0.27</u>	23.84	3.55	2.73
WyFormer-DFT	0.25	0.010	0.26	0.73	0.49	0.26	<u>23.56</u>	<u>3.49</u>	2.74
ADiT	0.51	0.002	0.27	0.71	0.03	0.24	14.18	3.52	2.78
Crystal-GFN	0.55	0.121	1.70	0.75	0.29	0.32	32.06	3.48	2.34
Crystalformer	0.26	<u>0.003</u>	0.74	0.71	0.35	0.31	18.59	3.78	2.82
DiffCSP	0.45	0.008	0.11	0.73	0.13	0.24	14.42	3.29	<u>2.60</u>
DiffCSP++	0.21	<u>0.003</u>	0.08	0.72	0.51	<u>0.27</u>	20.84	3.56	2.77
LLaMat2	0.32	0.002	0.16	0.71	0.25	0.24	9.86	3.94	2.86
LLaMat3	0.34	0.005	5.14	0.70	0.10	0.30	13.15	3.83	2.81
SymmCD	<u>0.22</u>	0.006	0.96	0.73	0.49	0.27	19.45	3.54	2.74
Alexandria	0.09	0.000	0.01	0.71	0.50	0.25	13.50	4.02	3.33
OQMD	0.27	0.004	0.14	0.63	0.48	0.24	12.57	4.22	3.29
AFLOW	0.28	0.004	0.12	0.69	0.46	0.25	10.07	4.12	3.22
MP-Exp	0.33	0.021	0.69	0.72	0.71	0.28	54.79	2.78	2.27

Table 9 Model evaluation metrics calculated using **uma** with **LeMat-Bulk** as the reference set **after relaxation**: Validity, Energy, Stability, and Metastability. All models provide 2500 structures except Crystalformer (1000 structures). Submissions are organized into four categories separated by horizontal lines: (1) models that incorporate relaxation as part of their generation pipeline (MatterGen, PLaID++, WyFormer, WyFormer-DFT), (2) other generative models, (3) samples from real-world datasets that contribute to LeMat-Bulk (Alexandria, OQMD) as baselines, and (4) samples from the AFLOW dataset and MP-Exp structures. Best values are shown in **bold**, second-best values are underlined. Arrows indicate whether higher ($\uparrow$) or lower ($\downarrow$) values are better.

Model	Validity				Unique% $\uparrow$	Novel% $\uparrow$	Energy-based			Stability			Metastability		
	Valid% $\uparrow$	CN% $\uparrow$	MinDist% $\uparrow$	PhysPlau% $\uparrow$			FormE (Std) $\downarrow$	E_{bulk} (Std) $\downarrow$	RMSD (Std) $\downarrow$	Stable% $\uparrow$	U-Stable% $\uparrow$	SUN% $\uparrow$	Metastable% $\uparrow$	U-Meta% $\uparrow$	MSUN% $\uparrow$
MatterGen	95.9	95.9	100.0	100.0	95.3	<u>70.4</u>	-0.74 ± 0.80	0.14 ± 0.14	0.00 ± 0.02	4.5 (113)	4.4 (111)	0.8 (20)	39.5 (987)	39.0 (976)	21.7 (543)
PLaID++	99.7	99.8	100.0	100.0	80.8	24.7	-0.49 ± 0.47	0.08 ± 0.15	0.00 ± 0.01	16.8 (404)	11.3 (272)	1.3 (32)	59.4 (1430)	46.7 (1123)	8.2 (198)
WyFormer	95.2	95.2	100.0	100.0	94.7	66.1	-0.68 ± 0.89	0.25 ± 0.28	0.01 ± 0.05	4.0 (101)	3.9 (97)	0.5 (12)	27.5 (687)	27.2 (679)	10.4 (260)
WyFormer-DFT	95.2	95.2	100.0	100.0	95.0	66.3	-0.71 ± 0.89	0.22 ± 0.25	0.00 ± 0.02	3.8 (94)	3.6 (90)	0.4 (10)	27.7 (691)	27.6 (689)	10.3 (257)
ADiT	<u>99.0</u>	<u>99.0</u>	<u>99.9</u>	100.0	<u>95.2</u>	24.3	-1.04 ± 0.90	0.06 ± 0.09	0.00 ± 0.01	17.7 (441)	16.5 (413)	0.6 (16)	62.6 (1562)	60.3 (1506)	10.0 (250)
Crystal-GFN	56.1	58.2	99.3	97.0	56.1	56.1	-2.54 ± 2.08	0.86 ± 1.74	0.12 ± 0.14	0.0 (0)	0.0 (0)	0.0 (0)	1.0 (24)	1.0 (24)	1.0 (24)
Crystalformer	94.7	96.3	98.4	100.0	94.2	49.4	-0.73 ± 0.86	0.17 ± 0.24	0.01 ± 0.08	10.0 (96)	9.9 (95)	0.8 (8)	41.2 (396)	40.9 (393)	10.2 (98)
DiffCSP	96.2	96.2	100.0	100.0	94.8	65.7	-0.76 ± 0.82	0.14 ± 0.15	0.01 ± 0.05	5.9 (148)	5.3 (133)	0.8 (20)	40.0 (1000)	39.5 (988)	19.4 (484)
DiffCSP++	95.5	95.5	100.0	100.0	<u>95.2</u>	59.8	-0.72 ± 0.86	0.20 ± 0.21	0.01 ± 0.04	5.3 (132)	5.2 (131)	0.5 (13)	31.6 (789)	31.4 (785)	10.1 (253)
LLaMat2	97.0	97.0	<u>99.9</u>	100.0	93.8	38.2	-0.77 ± 0.85	0.13 ± 0.19	0.00 ± 0.04	11.3 (282)	10.6 (266)	1.0 (24)	48.1 (1202)	45.8 (1145)	9.1 (227)
LLaMat3	83.7	84.6	98.6	100.0	77.4	83.4	-0.38 ± 0.71	0.27 ± 0.26	0.02 ± 0.06	0.8 (17)	0.8 (17)	<u>0.2 (4)</u>	17.6 (366)	17.5 (364)	14.5 (302)
SymmCD	95.6	95.9	99.7	100.0	<u>95.2</u>	60.7	-0.68 ± 0.84	0.23 ± 0.25	0.01 ± 0.05	5.6 (140)	5.5 (138)	0.6 (15)	30.6 (764)	30.3 (758)	9.4 (236)
Alexandria	94.3	94.3	100.0	100.0	94.3	4.6	-0.18 ± 0.95	0.43 ± 0.60	0.00 ± 0.02	1.7 (43)	1.7 (43)	0.0 (1)	28.4 (702)	28.4 (702)	0.8 (20)
OQMD	97.3	97.3	100.0	100.0	96.9	4.3	-0.22 ± 0.90	0.38 ± 0.45	0.00 ± 0.00	6.4 (158)	6.3 (156)	0.2 (4)	29.4 (730)	29.3 (729)	0.9 (22)
AFLOW	91.4	91.4	100.0	100.0	87.1	21.6	-0.18 ± 0.87	0.34 ± 0.42	0.00 ± 0.00	12.3 (307)	11.3 (282)	0.0 (0)	28.3 (708)	25.6 (639)	0.8 (21)
MP-Exp	98.0	98.0	100.0	100.0	85.9	8.0	-1.03 ± 0.89	0.07 ± 0.18	0.00 ± 0.02	25.3 (633)	22.4 (560)	1.0 (24)	55.4 (1386)	48.6 (1214)	2.0 (51)

Table 10 Model evaluation metrics calculated using **uma** with **LeMat-Bulk** as the reference set **after relaxation**: Distribution, Diversity, and HHI. All models provide 2500 structures except Crystalformer (1000 structures). Submissions are organized into four categories separated by horizontal lines: (1) models that incorporate relaxation as part of their generation pipeline (MatterGen, PLaID++, WyFormer, WyFormer-DFT), (2) other generative models, (3) samples from real-world datasets that contribute to LeMat-Bulk (Alexandria, OQMD) as baselines, and (4) samples from the AFLOW dataset and MP-Exp structures. Best values are shown in **bold**, second-best values are underlined. Arrows indicate whether higher ($\uparrow$) or lower ($\downarrow$) values are better.

Model	Distribution			Diversity				HHI	
	JS $\downarrow$	MMD $\downarrow$	FID $\downarrow$	ElemDiv $\uparrow$	SGDiv $\uparrow$	SizeDiv $\uparrow$	SiteDiv $\uparrow$	Prod $\downarrow$	Res $\downarrow$
MatterGen	0.35	0.006	0.51	0.64	0.32	0.24	12.81	3.51	2.72
PLaID++	0.43	0.029	1.02	0.68	0.30	0.21	6.01	5.23	3.36
WyFormer	<u>0.28</u>	0.008	<u>0.22</u>	0.74	0.51	<u>0.27</u>	24.27	3.52	2.71
WyFormer-DFT	0.29	0.010	0.26	<u>0.73</u>	0.49	<u>0.27</u>	<u>23.57</u>	<u>3.49</u>	2.74
ADiT	0.36	0.003	1.03	0.71	0.40	0.24	14.50	3.41	2.70
Crystal-GFN	0.64	0.092	2.94	0.74	0.12	0.32	31.50	3.45	2.33
Crystalformer	0.29	0.006	0.43	0.71	0.39	0.31	21.57	3.56	<u>2.69</u>
DiffCSP	0.35	0.009	0.30	<u>0.73</u>	0.34	0.25	14.42	3.27	2.59
DiffCSP++	0.27	0.004	0.43	0.72	<u>0.50</u>	0.27	20.90	3.55	2.77
LLaMat2	<u>0.28</u>	0.002	0.77	0.71	0.38	0.24	11.91	3.84	2.79
LLaMat3	0.43	<u>0.003</u>	0.15	0.72	0.17	0.27	28.26	3.48	2.93
SymmCD	0.28	0.008	0.28	<u>0.73</u>	0.50	0.27	20.87	3.51	2.68
Alexandria	0.14	0.000	0.01	0.71	0.50	0.26	13.50	4.02	3.33
OQMD	0.28	0.003	0.16	0.63	0.46	0.24	12.57	4.22	3.29
AFLOW	0.29	0.004	0.13	0.69	0.50	0.25	10.07	4.12	3.22
MP-Exp	0.35	0.022	0.74	0.72	0.68	0.28	54.79	2.78	2.27

C.3 Using a Self-Consistent MLIP-Based Convex Hull

A key challenge in evaluating generated structures with Machine-Learned Interatomic Potentials (MLIPs) is the accurate calculation of the energy above the convex hull (E_{hull}). Mixing total energies from an MLIP with a convex hull built from reference DFT calculations introduces inconsistencies: the MLIP and DFT methods operate on different energy references and possess distinct systematic errors. Since stability thresholds are tight (often a few tens of meV/atom), even small systematic shifts can change a structure’s classification from stable to unstable.

To address this, our benchmark constructs the convex hull using the same MLIP that evaluates the candidate structures. This self-consistent approach allows systematic errors to cancel, yielding more reliable E_{hull} estimates.

Empirical validation. **Figures 4** and **11** compare F1-scores for stability prediction under the two approaches. On both LeMat-Bulk and MP-20, the self-consistent hull (panel b) consistently outperforms the mixed-reference DFT hull (panel a), with the largest improvements at the strictest threshold (< 0.01 eV/atom). For example, on LeMat-Bulk the ORB F1-score improves from 0.66 to 0.81.

Figure 12 reveals why. The self-consistent approach (top row) yields systematically lower MAE and higher correlation with DFT ground truth than the mixed-reference approach (bottom row). For ORB, MAE decreases from 0.012 to 0.008 eV/atom; for UMA, from 0.012 to 0.007 eV/atom.

Ensemble composition. Although MACE-MP exhibits higher variance and lower correlation with DFT ($r \approx 0.7$) than ORB and UMA ($r > 0.9$), we retain it in the ensemble to ensure model diversity. ORB and UMA were both trained on OMat24 in addition to Materials Project, making their predictions partially correlated; relying on them alone risks collapsing the ensemble toward a single effective model. MACE-MP, trained exclusively on Materials Project, provides an independent perspective that guards against this collapse. On the self-consistent hull, systematic biases cancel within each model, and the remaining variance-type errors are effectively reduced by averaging across architectures with diverse training data. This design choice reflects a deliberate trade-off: accepting slightly higher individual-model error in exchange for more robust and diverse ensemble estimates.




Figure 11 F1-score for stability prediction on MP-20. Format matches [Figure 4](#).




Figure 12 Correlation between MLIP-predicted and DFT-calculated E_{hull} on LeMat-Bulk. *Top row (Own Hull)*: Self-consistent approach yields high correlation ($r > 0.94$ for ORB and UMA), low error ($\text{MAE} \leq 0.008$ eV/atom), and tight clustering around the diagonal. The ensemble (rightmost) achieves intermediate performance ($\text{MAE} = 0.011$ eV/atom). *Bottom row (DFT Hull)*: Mixed-reference approach shows degraded correlation and higher MAE. Notably, the ensemble ($\text{MAE} = 0.035$ eV/atom) performs worse than all individual models, illustrating how averaging heterogeneous biases can compound error.


C.4 Benchmark Results with MP-20 as the Reference Database

We visualize S.U.N. and M.S.U.N. structures from various generative models using t-SNE, a dimensionality reduction method in Figure 13. Across both plots, most models produce structures that are spread out broadly throughout the embedding space rather than collapsing into small, isolated regions. This suggests that current crystal generators are relatively capable of proposing chemically diverse candidates and are not restricted to narrow pockets of the MP-20 distribution. Although there is some overlap, different models still preferentially populate distinct regions of the embedding manifold. This complementary coverage supports the idea that combining multiple generative models within a discovery pipeline could increase the overall diversity of viable candidates.


One notable pattern appears in Figure 13b, where Crystal-GFN produces a comparatively tight, well-defined cluster. Despite having the second highest S.U.N. rate, this indicates the model may be biased toward a narrower slice of chemically stable structures relative to other methods. In contrast, models like PLaID++, WyFormer, and MatterGen reflect broader exploration of the structural landscape.

We report model efficiency metrics in Figure 14 and Figure 15. The Pareto frontier in inference-time space is again spanned by WyFormer, Crystalformer, and PLaID++, which interpolate between low-latency, moderate-quality generation and slower, higher S.U.N.% regimes. WyFormer remains at the extreme low-latency corner, whereas PLaID++ attains the highest S.U.N.% at substantially increased runtime, with MatterGen, DiffCSP, and other approaches lying close to but generally below this curve.

A similar picture appears when plotting S.U.N.% versus parameter count. WyFormer and Crystalformer achieve competitive S.U.N.% with relatively compact models, and DiffCSP and MatterGen extend this frontier to moderately larger parameter budgets. PLaID++ pushes to the highest S.U.N.% but does so with a substantially larger model, while SymmCD, ADiT, and the LLaMat variants sit well below the frontier—using one to two orders of magnitude more parameters than the most efficient models yet achieving comparable or worse S.U.N.%. Since PLaID++ directly optimizes for S.U.N. with respect to MP-20, it is unsurprising that it has the highest S.U.N. rate. At the same time, its performance highlights the promise of reinforcement learning-based optimization which could be applied to smaller models to improve material discovery rates across a range of model sizes.




(a) t-SNE visualization of ORB graph embeddings for M.S.U.N. materials.




(b) t-SNE visualization of ORB graph embeddings for S.U.N. materials.

Figure 13 Two-dimensional t-SNE visualizations (perplexity = 30) of ORB graph embeddings of generative model structures computed with MP-20 as the reference step. The visualizations demonstrate how most models generate a wide diversity of novel crystal structures which cover a breadth of chemical space.




(a) Inference time vs MSUN%




(b) Model parameters vs MSUN%

Figure 14 Performance-efficiency trade-offs for crystal structure generation models evaluated on MP-20 measured using MSUN. (a) Inference time (log scale) versus MSUN% shows the computational cost required to achieve MSUN structures. Models on the Pareto frontier (connected by blue line) represent optimal time-performance trade-offs. (b) Model parameter count (log scale) versus MSUN% illustrates the relationship between model size and generation quality. WyFormer, Crystalformer, DiffCSP, and MatterGen form the Pareto frontier.



(a) Inference time vs SUN%



(b) Model parameters vs SUN%

Figure 15 Performance-efficiency trade-offs for crystal structure generation models evaluated on MP-20 measured using SUN. (a) Inference time versus S.U.N.% highlights the computational budget required to produce stable and unique structures. The Pareto frontier shows that WyFormer, Crystalformer, and PLaID++ jointly define the best speed-quality trade-offs, with PLaID++ achieving the highest S.U.N.% at a moderate additional cost. (b) Model size versus S.U.N.% illustrates how parameter count correlates with generation quality. With WyFormer, CrystalFormer, DiffCSP and PLaID++ on the pareto frontier, this demonstrates that larger models can achieve better performance when properly designed.

Table 11 Model evaluation metrics calculated using **MLIP ensemble** with **MP-20** as the reference set: Validity, Energy, Stability, and Metastability. All models provide 2500 structures except Crystalformer (1000 structures). Submissions are organized into four categories separated by horizontal lines: (1) models that incorporate relaxation as part of their generation pipeline (MatterGen, PLaID++, WyFormer, WyFormer-DFT), (2) other generative models, (3) samples from real-world datasets that contribute to LeMat-Bulk (Alexandria, OQMD) as baselines, and (4) samples from the AFLOW dataset and Materials Project Experimental structures. Best values are shown in **bold**, second-best values are underlined. Arrows indicate whether higher ($\uparrow$) or lower ($\downarrow$) values are better.

Model	Validity				Unique% $\uparrow$	Novel% $\uparrow$	Energy-based			Stability			Metastability		
	Valid% $\uparrow$	CN% $\uparrow$	MinDist% $\uparrow$	PhysPlau% $\uparrow$			FormE (Std) $\downarrow$	E_{hull} (Std) $\downarrow$	RMSD (Std) $\downarrow$	Stable% $\uparrow$	U-Stable% $\uparrow$	SUN% $\uparrow$	Metastable% $\uparrow$	U-Meta% $\uparrow$	MSUN% $\uparrow$
MatterGen	95.7	95.9	99.8	100.0	<u>95.1</u>	83.1	-0.70 $\pm$ 0.79	0.16 $\pm$ 0.18	0.39 $\pm$ 0.50	6.0 (150)	6.0 (150)	3.6 (91)	34.6 (865)	<u>34.1</u> (853)	24.7 (618)
PLaID++	<u>96.0</u>	96.0	96.2	96.2	77.8	63.9	-0.50 $\pm$ 0.44	0.06 $\pm$ 0.17	0.13 $\pm$ 0.29	31.0 (776)	21.1 (528)	11.6 (289)	45.1 (1128)	38.0 (933)	<u>24.2</u> (605)
WyFormer	93.4	95.2	98.2	100.0	93.0	80.4	-0.43 $\pm$ 0.95	0.47 $\pm$ 0.52	0.81 $\pm$ 0.97	2.8 (71)	2.8 (70)	2.1 (52)	16.1 (405)	16.0 (400)	6.6 (166)
WyFormer-DFT	95.2	95.2	100.0	100.0	95.0	83.3	-0.67 $\pm$ 0.91	0.24 $\pm$ 0.38	0.42 $\pm$ 0.59	8.2 (205)	8.0 (201)	4.0 (100)	24.4 (610)	24.4 (609)	16.8 (419)
ADiT	90.6	99.0	91.5	100.0	87.8	31.3	<u>-0.73 $\pm$ 0.93</u>	0.32 $\pm$ 0.45	0.38 $\pm$ 0.40	1.0 (25)	1.0 (25)	0.6 (15)	<u>39.8</u> (994)	38.0 (951)	3.1 (78)
Crystal-GFN	51.7	58.3	91.1	96.3	51.7	51.7	-1.29 $\pm$ 2.65	1.66 $\pm$ 2.51	1.86 $\pm$ 0.90	5.6 (141)	5.6 (141)	<u>5.6</u> (141)	0.9 (22)	0.9 (22)	0.9 (22)
Crystalformer	69.9	96.3	72.9	<u>99.9</u>	69.4	44.2	-0.17 $\pm$ 1.48	0.68 $\pm$ 1.34	0.66 $\pm$ 1.03	5.2 (52)	5.2 (52)	3.3 (33)	27.8 (278)	27.5 (275)	9.4 (94)
DiffCSP	95.7	96.2	99.5	100.0	94.8	81.9	-0.64 $\pm$ 0.96	0.25 $\pm$ 0.56	0.59 $\pm$ 0.83	7.2 (179)	7.0 (176)	4.1 (102)	28.8 (720)	28.1 (702)	18.3 (457)
DiffCSP++	95.3	95.5	99.7	100.0	<u>95.1</u>	81.1	-0.52 $\pm$ 0.87	0.39 $\pm$ 0.44	0.70 $\pm$ 0.83	4.8 (121)	4.8 (121)	3.4 (85)	26.3 (657)	26.1 (652)	14.1 (353)
LLaMat2	84.4	<u>97.0</u>	87.1	99.6	81.4	53.0	-0.47 $\pm$ 1.01	0.42 $\pm$ 0.71	0.54 $\pm$ 0.71	3.9 (97)	3.8 (94)	2.9 (73)	35.1 (878)	32.9 (823)	10.6 (266)
LLaMat3	15.4	82.5	16.8	84.2	15.2	13.8	0.74 $\pm$ 2.69	1.68 $\pm$ 2.48	1.01 $\pm$ 1.10	0.5 (12)	0.5 (12)	0.4 (11)	2.1 (52)	2.0 (50)	0.8 (21)
SymmCD	73.4	95.9	76.4	100.0	73.0	61.3	-0.02 $\pm$ 1.22	0.85 $\pm$ 1.07	0.87 $\pm$ 1.10	4.3 (108)	4.3 (107)	2.2 (54)	17.8 (440)	17.6 (440)	9.0 (225)
Alexandria	93.4	93.4	99.0	99.0	93.4	<u>92.4</u>	-0.18 $\pm$ 0.94	0.41 $\pm$ 0.62	0.15 $\pm$ 0.36	10.7 (268)	10.7 (268)	10.5 (263)	24.6 (616)	24.6 (616)	24.0 (599)
OQMD	96.8	96.8	99.4	99.5	96.4	94.2	-0.23 $\pm$ 0.88	0.36 $\pm$ 0.47	0.14 $\pm$ 0.28	<u>17.5</u> (438)	<u>17.5</u> (437)	16.6 (415)	19.9 (497)	19.8 (495)	18.2 (456)
AFLOW	91.4	91.4	100.0	100.0	87.1	57.0	-0.19 $\pm$ 0.86	0.33 $\pm$ 0.42	0.12 $\pm$ 0.44	16.0 (401)	15.0 (376)	3.0 (74)	25.2 (631)	22.8 (571)	4.2 (105)
Experimental	98.0	98.0	100.0	100.0	85.6	58.6	-1.02 $\pm$ 0.97	0.08 $\pm$ 0.46	0.20 $\pm$ 0.36	40.5 (1015)	35.4 (885)	25.0 (624)	41.5 (1038)	36.6 (914)	13.0 (326)

Table 12 Model evaluation metrics calculated using **MLIP ensemble** with **MP-20** as the reference set: Distribution, Diversity, and HHI. All models provide 2500 structures except Crystalformer (1000 structures). Submissions are organized into four categories separated by horizontal lines: (1) models that incorporate relaxation as part of their generation pipeline (MatterGen, PLaID++, WyFormer, WyFormer-DFT), (2) other generative models, (3) samples from real-world datasets that contribute to LeMat-Bulk (Alexandria, OQMD) as baselines, and (4) samples from the AFLOW dataset and MP-Exp structures. Best values are shown in **bold**, second-best values are underlined. Arrows indicate whether higher ($\uparrow$) or lower ($\downarrow$) values are better.

Model	Distribution			Diversity				HHI	
	JS $\downarrow$	MMD $\downarrow$	FID $\downarrow$	ElemDiv $\uparrow$	SGDiv $\uparrow$	SizeDiv $\uparrow$	SiteDiv $\uparrow$	Prod $\downarrow$	Res $\downarrow$
MatterGen	0.43	0.000	14.38	0.64	0.14	0.21	12.78	3.52	2.72
PLaID++	0.48	0.049	15.48	0.68	0.26	0.21	6.01	5.23	3.36
WyFormer	0.42	<u>0.001</u>	15.03	<u>0.74</u>	<u>0.50</u>	0.27	<u>23.84</u>	3.55	2.73
WyFormer-DFT	0.42	0.002	21.23	0.73	0.49	0.26	23.56	3.49	2.74
ADiT	0.43	<u>0.001</u>	<u>14.41</u>	0.71	0.03	0.24	14.18	3.52	2.78
Crystal-GFN	0.56	0.098	50.20	0.75	0.29	0.32	32.06	<u>3.48</u>	2.34
Crystalformer	<u>0.41</u>	0.002	16.60	0.71	0.35	<u>0.31</u>	18.59	3.78	2.82
DiffCSP	<u>0.44</u>	0.002	14.73	0.73	0.13	0.24	14.42	3.29	<u>2.60</u>
DiffCSP++	0.42	0.003	14.97	0.72	0.51	0.27	20.84	3.56	2.77
LLaMat2	<u>0.41</u>	0.003	14.91	0.71	0.25	0.24	9.86	3.94	2.86
LLaMat3	<u>0.45</u>	0.012	27.67	0.70	0.10	0.30	13.15	3.83	2.81
SymmCD	<u>0.41</u>	<u>0.001</u>	16.33	0.73	0.49	0.27	19.45	3.54	2.74
Alexandria	0.38	0.004	14.82	0.71	<u>0.50</u>	0.25	13.50	4.02	3.33
OQMD	0.43	0.014	15.32	0.63	0.48	0.24	12.57	4.22	3.29
AFLOW	0.44	0.003	15.34	0.69	0.46	0.25	10.07	4.12	3.22
MP-Exp	0.49	0.011	14.73	0.72	0.71	0.28	54.79	2.78	2.27

Table 13 Model evaluation metrics calculated using **uma** with **MP-20** as the reference set: Validity, Energy, Stability, and Metastability. All models provide 2500 structures except Crystalformer (1000 structures). Submissions are organized into four categories separated by horizontal lines: (1) models that incorporate relaxation as part of their generation pipeline (MatterGen, PLaID++, WyFormer, WyFormer-DFT), (2) other generative models, (3) samples from real-world datasets that contribute to LeMat-Bulk (Alexandria, OQMD) as baselines, and (4) samples from the AFLOW dataset and MP-Exp structures. Best values are shown in **bold**, second-best values are underlined. Arrows indicate whether higher ($\uparrow$) or lower ($\downarrow$) values are better.

Model	Validity				Unique% $\uparrow$	Novel% $\uparrow$	Energy-based			Stability			Metastability		
	Valid% $\uparrow$	CN% $\uparrow$	MinDist% $\uparrow$	PhysPlau% $\uparrow$			FormE (Std) $\downarrow$	E_{hull} (Std) $\downarrow$	RMSD (Std) $\downarrow$	Stable% $\uparrow$	U-Stable% $\uparrow$	SUN% $\uparrow$	Metastable% $\uparrow$	U-Meta% $\uparrow$	MSUN% $\uparrow$
MatterGen	95.7	95.9	99.8	100.0	<u>95.1</u>	83.1	-0.69 $\pm$ 0.80	0.17 $\pm$ 0.19	0.40 $\pm$ 0.60	5.4 (136)	5.4 (135)	3.5 (88)	33.4 (835)	33.0 (824)	23.1 (578)
PLaID++	<u>96.0</u>	96.0	96.2	96.2	77.8	63.9	-0.48 $\pm$ 0.48	0.07 $\pm$ 0.17	0.12 $\pm$ 0.32	29.6 (741)	20.6 (514)	12.0 (299)	45.6 (1141)	37.4 (935)	<u>23.1</u> (578)
WyFormer	93.4	95.2	98.2	100.0	93.0	80.4	-0.41 $\pm$ 0.97	0.49 $\pm$ 0.54	0.80 $\pm$ 1.05	2.4 (61)	2.4 (60)	1.9 (47)	16.2 (405)	16.0 (401)	6.7 (168)
WyFormer-DFT	95.2	95.2	100.0	100.0	95.0	83.3	-0.65 $\pm$ 1.07	0.26 $\pm$ 0.66	0.39 $\pm$ 0.60	7.5 (187)	7.4 (184)	3.8 (94)	24.4 (611)	24.4 (609)	16.4 (409)
ADiT	90.6	99.0	91.5	100.0	87.8	31.3	<u>-0.72 $\pm$ 0.95</u>	0.32 $\pm$ 0.45	0.38 $\pm$ 0.41	1.0 (26)	1.0 (26)	0.6 (16)	<u>38.2</u> (954)	36.6 (915)	3.0 (76)
Crystal-GFN	51.7	58.3	91.1	96.3	51.7	51.7	-0.56 $\pm$ 4.46	2.38 $\pm$ 4.28	1.86 $\pm$ 1.10	5.4 (134)	5.4 (134)	5.4 (134)	0.8 (19)	0.8 (19)	0.8 (19)
Crystalformer	69.9	96.3	72.9	99.9	69.4	44.2	-0.15 $\pm$ 1.48	0.69 $\pm$ 1.33	0.68 $\pm$ 1.24	5.1 (51)	5.1 (51)	3.2 (32)	27.9 (279)	27.8 (278)	9.5 (95)
DiffCSP	95.7	96.2	99.5	100.0	94.8	81.9	-0.62 $\pm$ 0.98	0.26 $\pm$ 0.58	0.59 $\pm$ 0.65	6.4 (160)	6.2 (155)	3.8 (95)	28.8 (719)	28.2 (704)	17.8 (445)
DiffCSP++	95.3	95.5	99.7	100.0	<u>95.1</u>	81.1	-0.50 $\pm$ 0.89	0.40 $\pm$ 0.45	0.69 $\pm$ 0.88	4.7 (118)	4.7 (118)	3.2 (79)	25.6 (640)	25.4 (635)	13.6 (341)
LLaMat2	84.4	<u>97.0</u>	87.1	99.6	81.4	53.0	-0.45 $\pm$ 1.03	0.43 $\pm$ 0.72	0.54 $\pm$ 0.75	4.6 (116)	4.5 (112)	3.2 (80)	33.3 (832)	31.3 (782)	9.7 (243)
LLaMat3	15.4	82.5	16.8	84.2	15.2	13.8	0.79 $\pm$ 2.73	1.72 $\pm$ 2.52	1.01 $\pm$ 1.26	0.4 (10)	0.4 (10)	0.4 (9)	2.2 (55)	2.2 (54)	1.0 (24)
SymmCD	73.4	95.9	76.4	100.0	73.0	61.3	-0.00 $\pm$ 1.23	0.87 $\pm$ 1.08	0.85 $\pm$ 1.13	3.7 (92)	3.6 (91)	2.3 (57)	18.0 (456)	18.0 (451)	8.8 (219)
Alexandria	93.4	93.4	99.0	99.0	93.4	<u>92.4</u>	-0.16 $\pm$ 0.97	0.42 $\pm$ 0.64	<u>0.12 $\pm$ 0.40</u>	9.9 (247)	9.9 (247)	9.8 (245)	24.5 (612)	24.5 (612)	23.7 (592)
OQMD	96.8	96.8	99.4	99.5	96.4	94.2	-0.21 $\pm$ 0.89	0.37 $\pm$ 0.47	0.13 $\pm$ 0.29	<u>17.2</u> (429)	<u>17.1</u> (428)	16.3 (407)	19.6 (491)	19.6 (489)	18.0 (449)
AFLOW	91.4	91.4	100.0	100.0	87.1	57.0	-0.18 $\pm$ 0.86	0.33 $\pm$ 0.42	0.11 $\pm$ 0.40	12.9 (322)	12.1 (302)	2.6 (66)	28.0 (699)	25.3 (632)	4.5 (112)
MP-Exp	98.0	98.0	100.0	100.0	85.6	58.6	-0.99 $\pm$ 1.36	0.08 $\pm$ 1.06	0.19 $\pm$ 0.39	35.0 (874)	30.5 (763)	23.8 (594)	46.7 (1168)	41.1 (1027)	14.0 (351)

Table 14 Model evaluation metrics calculated using **uma** with **MP-20** as the reference set: Distribution, Diversity, and HHI. All models provide 2500 structures except Crystalformer (1000 structures). Submissions are organized into four categories separated by horizontal lines: (1) models that incorporate relaxation as part of their generation pipeline (MatterGen, PLaID++, WyFormer, WyFormer-DFT), (2) other generative models, (3) samples from real-world datasets that contribute to LeMat-Bulk (Alexandria, OQMD) as baselines, and (4) samples from the AFLOW dataset and MP-Exp structures. Best values are shown in **bold**, second-best values are underlined. Arrows indicate whether higher ($\uparrow$) or lower ($\downarrow$) values are better.

Model	Distribution			Diversity				HHI	
	JS $\downarrow$	MMD $\downarrow$	FID $\downarrow$	ElemDiv $\uparrow$	SGDiv $\uparrow$	SizeDiv $\uparrow$	SiteDiv $\uparrow$	Prod $\downarrow$	Res $\downarrow$
MatterGen	0.43	0.000	0.18	0.64	0.14	0.24	12.78	3.52	2.72
PLaID++	0.48	0.049	0.30	0.68	0.26	0.21	6.01	5.23	3.36
WyFormer	0.42	0.001	1.24	<u>0.74</u>	<u>0.50</u>	0.27	<u>23.84</u>	3.55	2.73
WyFormer-DFT	0.42	0.002	<u>0.18</u>	0.73	0.49	0.26	23.56	3.49	2.74
ADiT	0.43	0.001	0.37	0.71	0.03	0.24	14.18	3.52	2.78
Crystal-GFN	0.56	0.098	1.70	0.75	0.29	0.32	32.06	<u>3.48</u>	2.34
Crystalformer	<u>0.41</u>	0.002	2.24	0.71	0.35	<u>0.31</u>	18.59	3.78	2.82
DiffCSP	0.44	0.002	0.60	0.73	0.13	0.24	14.42	3.29	<u>2.60</u>
DiffCSP++	0.42	0.003	0.09	0.72	0.51	0.27	20.84	3.56	2.77
LLaMat2	0.41	0.003	0.66	0.71	0.25	0.24	9.86	3.94	2.86
LLaMat3	0.45	0.012	9.31	0.70	0.10	0.30	13.15	3.83	2.81
SymmCD	0.41	<u>0.001</u>	3.34	0.73	0.49	0.27	19.45	3.54	2.74
Alexandria	0.38	0.004	0.97	0.71	0.50	0.25	13.50	4.02	3.33
OQMD	0.43	0.014	0.55	0.63	0.48	0.24	12.57	4.22	3.29
AFLOW	0.44	0.003	0.55	0.69	0.46	0.25	10.07	4.12	3.22
MP-Exp	0.49	0.011	0.16	0.72	0.71	0.28	54.79	2.78	2.27

Table 15 Model evaluation metrics calculated using **uma** with **MP-20** as the reference set **after relaxation**: Validity, Energy, Stability, and Metastability. All models provide 2500 structures except Crystalformer (1000 structures). Submissions are organized into four categories separated by horizontal lines: (1) models that incorporate relaxation as part of their generation pipeline (MatterGen, PLaID++, WyFormer, WyFormer-DFT), (2) other generative models, (3) samples from real-world datasets that contribute to LeMat-Bulk (Alexandria, OQMD) as baselines, and (4) samples from the AFLOW dataset and MP-Exp structures. Best values are shown in **bold**, second-best values are underlined. Arrows indicate whether higher ($\uparrow$) or lower ($\downarrow$) values are better.

Model	Validity				Unique% $\uparrow$	Novel% $\uparrow$	Energy-based			Stability			Metastability		
	Valid% $\uparrow$	CN% $\uparrow$	MinDist% $\uparrow$	PhysPlau% $\uparrow$			FormE (Std) $\downarrow$	E _{hull} (Std) $\downarrow$	RMSD (Std) $\downarrow$	Stable% $\uparrow$	U-Stable% $\uparrow$	SUN% $\uparrow$	Metastable% $\uparrow$	U-Meta% $\uparrow$	MSUN% $\uparrow$
MatterGen	95.8	95.8	100.0	100.0	<u>95.3</u>	83.0	-0.74 $\pm$ 0.80	0.12 $\pm$ 0.14	0.00 $\pm$ 0.05	10.2 (255)	10.1 (252)	5.8 (145)	40.2 (1006)	39.8 (995)	31.9 (797)
PLaID++	96.0	96.0	96.2	96.2	77.7	63.8	-0.49 $\pm$ 0.47	0.06 $\pm$ 0.16	0.00 $\pm$ 0.01	34.3 (857)	23.4 (586)	12.8 (320)	42.3 (1057)	35.4 (886)	23.4 (585)
WyFormer	95.0	95.1	99.8	99.9	94.6	81.7	-0.67 $\pm$ 0.90	0.22 $\pm$ 0.32	0.00 $\pm$ 0.02	9.6 (239)	<u>9.3 (233)</u>	4.8 (120)	26.6 (664)	26.3 (658)	18.1 (452)
WyFormer-DFT	95.2	95.2	100.0	100.0	94.9	83.1	-0.71 $\pm$ 0.89	0.20 $\pm$ 0.28	0.04 $\pm$ 0.21	9.4 (236)	9.3 (232)	4.7 (118)	26.6 (666)	26.6 (664)	19.4 (486)
ADiT	98.9	99.0	99.9	100.0	95.1	31.9	-1.04 $\pm$ 0.90	0.04 $\pm$ 0.00	0.00 $\pm$ 0.00	33.9 (848)	32.2 (805)	5.6 (140)	50.0 (1251)	48.3 (1208)	13.4 (335)
Crystal-GFN	55.3	57.4	98.0	95.7	55.3	55.3	-2.54 $\pm$ 2.08	0.41 $\pm$ 1.90	0.72 $\pm$ 0.75	15.6 (389)	15.6 (389)	15.6 (389)	3.4 (84)	3.4 (84)	3.4 (84)
Crystalformer	92.0	93.8	95.6	97.2	91.5	64.2	-0.74 $\pm$ 0.85	0.15 $\pm$ 0.25	0.00 $\pm$ 0.02	17.9 (179)	17.6 (176)	6.9 (69)	34.7 (347)	34.6 (346)	18.3 (183)
DiffCSP	96.2	96.2	100.0	100.0	94.8	81.8	-0.76 $\pm$ 0.82	0.12 $\pm$ 0.18	0.07 $\pm$ 0.28	13.4 (336)	12.8 (321)	7.1 (178)	38.6 (965)	38.0 (951)	30.4 (761)
DiffCSP++	95.5	95.5	<u>100.0</u>	<u>100.0</u>	95.2	80.8	-0.72 $\pm$ 0.86	0.18 $\pm$ 0.28	0.04 $\pm$ 0.21	12.2 (304)	12.1 (302)	6.8 (171)	29.1 (728)	29.0 (725)	19.9 (498)
LLaMat2	<u>97.0</u>	<u>97.0</u>	99.9	100.0	93.8	63.2	-0.77 $\pm$ 0.85	0.11 $\pm$ 0.21	0.03 $\pm$ 0.20	23.9 (597)	22.6 (566)	8.5 (213)	39.2 (980)	37.6 (939)	20.4 (510)
LLaMat3	69.8	70.5	82.2	83.3	69.6	68.0	-0.38 $\pm$ 0.71	0.24 $\pm$ 0.27	0.17 $\pm$ 0.46	2.5 (63)	2.5 (62)	1.7 (43)	16.0 (399)	15.9 (398)	15.0 (374)
SymmCD	95.6	95.9	99.7	100.0	95.1	79.8	-0.67 $\pm$ 0.85	0.21 $\pm$ 0.27	0.00 $\pm$ 0.03	12.5 (313)	12.3 (308)	6.0 (151)	27.4 (686)	27.3 (683)	18.4 (461)
Alexandria	93.4	93.4	99.0	99.0	93.4	<u>92.4</u>	-0.18 $\pm$ 0.95	0.40 $\pm$ 0.61	0.00 $\pm$ 0.02	10.5 (263)	10.5 (263)	10.3 (258)	24.8 (621)	24.8 (621)	24.2 (604)
OQMD	96.8	96.8	99.5	99.5	96.4	94.3	-0.22 $\pm$ 0.90	0.35 $\pm$ 0.47	<u>0.00 $\pm$ 0.00</u>	18.6 (466)	18.5 (463)	17.3 (432)	19.9 (497)	19.9 (497)	18.6 (466)
AFLOW	91.4	91.4	100.0	100.0	87.1	56.9	-0.18 $\pm$ 0.87	0.32 $\pm$ 0.42	0.00 $\pm$ 0.00	20.0 (499)	18.5 (462)	3.2 (79)	22.0 (549)	19.7 (492)	4.3 (108)
MP-Exp	98.0	98.0	100.0	100.0	85.9	58.9	-1.03 $\pm$ 0.89	0.04 $\pm$ 0.19	0.00 $\pm$ 0.02	50.3 (1257)	43.8 (1096)	26.0 (651)	32.9 (822)	29.2 (730)	13.3 (332)




Figure 16 Generative models evaluation across all metric categories. Top row: MP-20 reference. Middle row: LeMat-Bulk pre-relaxation. Bottom row: LeMat-Bulk post-relaxation. Columns from left to right: SUN Metrics, MSUN Metrics, Energy & Stability, Distribution & Diversity. All rows use consistent scaling for direct comparison. Models include relaxation-based generators (MatterGen, PLaID++, WyFormer, WyFormer-DFT) and diffusion/autoregressive approaches.




Figure 17 Dataset baselines evaluation across all metric categories. Top row: MP-20 reference. Middle row: LeMat-Bulk pre-relaxation. Bottom row: LeMat-Bulk post-relaxation. Columns from left to right: SUN Metrics, MSUN Metrics, Energy & Stability, Distribution & Diversity. All rows use consistent scaling for direct comparison. Baselines include Alexandria, OQMD, AFLOW, and MP-Exp.

Table 16 Model evaluation metrics calculated using **uma** with **MP-20** as the reference set **after relaxation**: Distribution, Diversity, and HHI. All models provide 2500 structures except Crystalformer (1000 structures). Submissions are organized into four categories separated by horizontal lines: (1) models that incorporate relaxation as part of their generation pipeline (MatterGen, PLaID++, WyFormer, WyFormer-DFT), (2) other generative models, (3) samples from real-world datasets that contribute to LeMat-Bulk (Alexandria, OQMD) as baselines, and (4) samples from the AFLOW dataset and MP-Exp structures. Best values are shown in **bold**, second-best values are underlined. Arrows indicate whether higher ($\uparrow$) or lower ($\downarrow$) values are better.

Model	Distribution			Diversity				HHI	
	JS $\downarrow$	MMD $\downarrow$	FID $\downarrow$	ElemDiv $\uparrow$	SGDiv $\uparrow$	SizeDiv $\uparrow$	SiteDiv $\uparrow$	Prod $\downarrow$	Res $\downarrow$
MatterGen	0.43	0.000	0.10	0.64	0.30	0.24	12.80	3.51	2.72
PLaID++	0.48	0.049	0.31	0.68	0.28	0.21	6.01	5.23	3.36
WyFormer	0.42	0.002	0.35	0.74	<u>0.50</u>	0.27	24.27	3.52	2.71
WyFormer-DFT	0.42	0.002	0.30	0.73	0.49	0.27	23.57	3.49	2.74
ADiT	0.43	<u>0.000</u>	0.02	0.71	0.43	0.24	14.51	<u>3.41</u>	2.70
Crystal-GFN	0.56	0.066	1.32	<u>0.74</u>	0.12	0.32	31.50	3.45	2.33
Crystalformer	0.42	0.002	0.13	0.71	0.34	<u>0.31</u>	21.62	3.54	2.68
DiffCSP	0.44	0.003	0.26	0.73	0.34	0.25	14.42	3.27	<u>2.59</u>
DiffCSP++	0.42	0.001	0.09	0.72	0.50	0.27	20.90	3.55	2.77
LLaMat2	<u>0.42</u>	0.001	<u>0.05</u>	0.71	0.38	0.24	11.91	3.84	2.79
LLaMat3	0.44	0.007	0.68	0.72	0.17	0.27	<u>28.26</u>	3.48	2.93
SymmCD	0.42	0.002	0.27	0.73	0.48	0.27	20.84	3.51	2.68
Alexandria	0.38	0.004	0.91	0.71	0.50	0.26	13.50	4.02	3.33
OQMD	0.43	0.012	0.51	0.63	0.46	0.24	12.57	4.22	3.29
AFLOW	0.44	0.003	0.53	0.69	0.50	0.25	10.07	4.12	3.22
MP-Exp	0.49	0.013	0.12	0.72	0.68	0.28	54.79	2.78	2.27

D Environmental and Sustainability Considerations

The application of generative models to materials discovery presents significant opportunities for advancing environmental sustainability goals. As global challenges related to climate change, resource depletion, and environmental degradation intensify, the need for novel materials with reduced environmental footprints becomes increasingly urgent. Generative approaches can accelerate the discovery of sustainable alternatives by explicitly incorporating environmental criteria into the design process.

One promising direction involves the targeted generation of materials with reduced reliance on critical or environmentally problematic elements. By conditioning generative models on compositional constraints that exclude toxic, rare, or environmentally harmful elements, researchers can guide exploration toward more sustainable regions of chemical space. Similarly, models can be trained to prioritize earth-abundant elements and avoid those associated with problematic extraction practices or geopolitical supply risks.

Energy-related applications represent another frontier where generative models could significantly impact sustainability outcomes. The discovery of more efficient catalysts for renewable energy production, improved battery materials for energy storage, and novel photovoltaic materials could accelerate the transition away from fossil fuels. By specifically targeting properties relevant to these applications, generative models can focus computational and experimental resources on high-impact sustainability domains.

Life-cycle considerations present a more complex but equally important target for integration with generative approaches. Ideally, materials should be designed not only for performance but also for recyclability, biodegradability, or other end-of-life scenarios that minimize environmental impact. Incorporating such considerations into generative frameworks remains challenging due to the complex, multi-faceted nature of life-cycle assessment, but represents a crucial direction for future research.

The computational efficiency of generative processes themselves also warrants consideration from a sustainability perspective. As models grow in complexity and scale, their energy consumption and carbon footprint increase accordingly. Developing more efficient architectures, training procedures, and sampling approaches could reduce the environmental impact of the discovery process itself, aligning computational means with environmental ends. This consideration becomes particularly important as generative approaches scale to industrial applications and high-throughput discovery platforms.

The ultimate success of generative approaches in advancing sustainability will depend not only on technical

capabilities but also on intentional alignment with environmental objectives. By explicitly incorporating sustainability metrics into reward functions, objective functions, and evaluation criteria, the materials community can ensure that generative models contribute to addressing environmental challenges rather than merely accelerating traditional discovery paradigms without regard for sustainability implications.