Inverse Design of Novel Superconductors via Guided Diffusion

Pawan Prakash
Department of Physics
Quantum Theory Project
University of Florida

Jason B. Gibson
Quantum Theory Project
Dept. of Materials Science and Engineering
University of Florida

Zhongwei Li
Department of Physics
University of Florida

Gabriele Di Gianluca
Department of Physics
University of Florida

Juan Esquivel
Department of Physics
University of Florida

Eric Fuemmeler
Dept. of Aerospace Engineering and Mechanics
University of Minnesota

Benjamin Geisler
Department of Physics
Dept. of Materials Science and Engineering
University of Florida

Jung Soo Kim
Department of Physics
University of Florida

Adrian Roitberg
Quantum Theory Project
Department of Chemistry
University of Florida

Ellad B. Tadmor
Dept. of Aerospace Engineering and Mechanics
University of Minnesota

Mingjie Liu
Quantum Theory Project
Department of Chemistry
University of Florida

Stefano Martiniani
Center for Soft Matter Research
Simons Center for Computational Physical Chemistry
New York University

Gregory R. Stewart
Department of Physics
University of Florida

James J. Hamlin
Department of Physics
University of Florida

Peter J. Hirschfeld
Department of Physics
University of Florida

Richard G. Hennig
Dept. of Materials Science and Engineering
Quantum Theory Project
University of Florida

Abstract

The inverse design of materials with specific desired properties, such as high-temperature superconductivity, represents a formidable challenge in materials science due to the vastness of chemical and structural space. This work introduces a guided diffusion framework to accelerate the discovery of novel superconductors.

39th Conference on Neural Information Processing Systems (NeurIPS 2025) Workshop: AI4Mat.

Our approach begins with a foundation model trained on crystal structures from the Alexandria Database, which is then fine-tuned using a labeled dataset of 7,217 conventional superconductors to generate new structures conditioned on critical temperature, T_c. Employing classifier-free guidance, we generated 200,000 potential crystal structures. These candidates were subsequently subjected to a rigorous multi-stage computational screening workflow, utilizing machine learning models and density functional theory calculations to assess stability and electronic properties. Notably, our generative model demonstrated effective property-driven design by shifting the distribution of generated materials toward targeted T_c values within the training regime. This process successfully identified 773 promising superconducting candidates with predicted $T_c > 5K$. This end-to-end workflow, from generation to new candidate superconductors, illuminates a powerful pathway for materials discovery, demonstrating the significant potential of the AI-driven framework to accelerate discovery.

1 Introduction

The discovery of novel materials with desired properties remains a fundamental challenge: traditional routes—direct simulation or experimental synthesis and characterization—are costly, slow, and yield few successes. While first principles methods and machine-learning (ML) models have improved property prediction from known crystal structures, the inverse problem—designing materials for target properties—remains daunting due to the vastness of chemical and structural space.

Generative models, which have achieved remarkable success in domains such as images, text, and video, are now gaining traction in materials science. Recent advances in structure-generating models, such as diffusion-based frameworks Ho et al. [2020] (e.g., DiffCSP Jiao et al. [2023], MatterGen Zeni et al. [2024]), flow-based Lipman et al. [2023] models (e.g., FlowMM Miller et al. [2024]), autoregressive transformers Vaswani et al. [2017] (e.g. CrystalLLM Antunes et al. [2024], Matra-Genoa De Breuck et al. [2025]), and frameworks based on stochastic interpolants Albergo et al. [2023] that unify diffusion and flow (e.g., OMatG Höllmer et al. [2025]), highlight the potential of generative approaches for material discovery.

Superconductors, with their zero electrical resistance, hold transformative potential for technologies ranging from energy transmission and storage to high-field applications in medical imaging, particle accelerators, and materials processing Boeri et al. [2022], Larbalestier et al. [2001], Malozemoff et al. [2005], Flükiger [2012]. These wide-ranging applications make the discovery of new superconductors, particularly those with higher critical temperatures (T_c), a high-priority target for inverse materials design. In the dominant class of electron-phonon superconductors, T_c is governed by the coupling between electrons and lattice vibrations Bardeen et al. [1957], Eliashberg [1960]. Traditionally density functional theory (DFT) Giannozzi et al. [2020] coupled to Eliashberg theory Poncé et al. [2016] has been used to accurately predict the T_c of superconductors, albeit at too high a computational cost to be suitable for high-throughput screening. Machine-learning implementations of Eliashberg theory Xie et al. [2019, 2022a] and of electron–phonon coupling Gibson et al. [2025a,b] are potential routes to accelerate the prediction of superconducting properties and screening efforts, enabling high-throughput screening.

Generative AI methods have the capability to directly propose crystal structures biased toward superconductivity by conditioning on target properties, potentially increasing the success rate relative to screening-only workflows. Generative-AI efforts at superconductor discovery include Wines *et al.* Wines et al. [2023], who combined a crystal diffusion variational autoencoder (CDVAE Xie et al. [2022b]) with ALIGNN Choudhary and DeCost [2021] as the property predictor. That study was a valuable proof of concept, but ALIGNN was trained on a small dataset (1,058 structures), limiting generalization, and the generated structures were reported in the low-symmetry space group $P\bar{1}$, which may bias diversity and physical realizability. Dordevic *et al.* Yuan and Dordevic [2024] introduced SuperDiff, a diffusion-based model that uses ILVR Choi et al. [2021] to condition on reference compounds; while fast and the first to generate new superconductor families via conditioning, it operates at the composition level without explicit crystal structures, limiting its ability to capture atomic arrangements. Other notable advances in AI-driven materials-discovery include Pogue *et al.* Pogue et al. [2023], Wilfong *et al.* Wilfong et al. [2025], Hutcheon *et al.* Hutcheon et al. [2020],

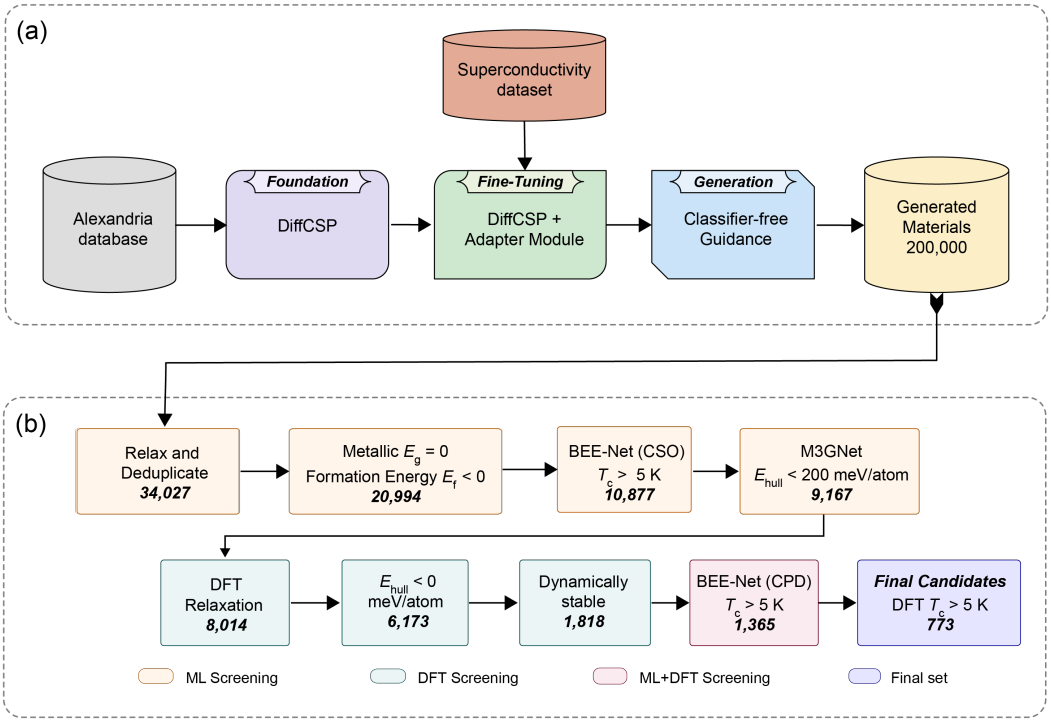

Figure 1: Workflow combining (a) guided diffusion and (b) multi-stage filtering to predict candidate superconductors. (a) Overview of the guided diffusion pipeline for superconductor discovery, starting from the Alexandria database of 2,086,767 crystal structures used here to train a DiffCSP-based foundation model capable of generating plausible crystals. An adapter module is implemented in the DiffCSP denoiser to fine-tune the foundation model using a dataset of 7,183 superconductors Cerqueira et al. [2023], Gibson et al. [2025b]. Using classifier free guidance, we use this AI framework to generate 200,000 crystal structures. (b) Overview of the multi-stage filtering process of the generated structures. We first pass the generated structures through initial ML relaxation and deduplication, followed by search for metallic, thermodynamically stable, ML predicted $T_c > 5$ K structures and $E_{hull} < 200$ meV/atom. We perform dynamic stability verification on the remaining candidates, and calculated the stable structures' T_c with BEE-NET Gibson et al. [2025b] using the PhDOS as embedding, selecting the $T_c > 5$ K. The electron-phonon spectral function of the final candidates is calculated using DFT, which is then used to obtain the final (DFT) T_c of the candidate structures.

Chen *et al.* Chen et al. [2021], and Griesemer *et al.* Griesemer et al. [2025], with additional context reviewed in Ref. Madika et al. [2025].

The field still lacks a scalable generative framework for superconductors that solves the inverse design problem under data scarcity conditions. Here, we build on recent advances in generative AI and present an enhanced DiffCSP framework that uses guided diffusion to generate crystal structures conditioned on target T_c values. To address the limited size of labeled superconductivity datasets, we decouple structural priors from property conditioning: we first pretrain DiffCSP on more than two million crystal structures from the Alexandria Database Schmidt et al. [2023a, 2022a], teaching the model to generate plausible crystals independent of any target property; we then fine-tune our model for T_c conditioning on a smaller, high-quality set of 7,217 conventional superconductors with calculated T_c Cerqueira et al. [2023] (Fig. 1a). This two-stage strategy leverages a massive corpus to learn structural validity while requiring only a specialized dataset to capture superconductivity-relevant correlations, thereby steering generation toward candidates that are both plausible and likely to be superconducting.

We apply this end-to-end workflow at scale by generating 200,000 candidate crystal structures (Fig. 1a) and filtering them through a multistage screening process that combines machine-learning models Chen et al. [2019], Chen and Ong [2022], Gibson et al. [2025b] with DFT Giannozzi et al. [2020]. The pipeline enforces metallicity (bandgap $E_{\mathrm{g}} = 0$), formation energy $E_{\mathrm{f}} < 0$, thermodynamic stability ($E_{\mathrm{hull}} < 200$ meV/atom computed via ML and DFT, where E_{hull} is the energy above the hull), and dynamical stability (phonons), and prioritizes superconducting propensity. This process yields 773 novel candidates with DFT-calculated $T_{\mathrm{c}} > 5$ K and $E_{\mathrm{hull}} < 200$ meV/atom (Fig. 1b).

2 Methods

Our methodology for discovering novel superconductors integrates two primary components: (1) a generative AI framework to propose novel crystal structures conditioned on superconducting properties; (2) a high-throughput computational screening workflow driven by machine learning models and ab initio calculations to assess the stability and promise of the generated candidates. The components of our discovery pipeline, illustrated in Figure 1, is designed to first generate a vast number of candidate structures and then systematically filter them to identify the most promising ones. The generative AI stage itself involves three core steps: training a foundation model, fine-tuning it for superconductivity, and guiding the generation process with classifier-free diffusion to create an initial ensemble of 200,000 structures (Fig. 1(a)). These candidates are then subjected to a rigorous, multi-stage filtering workflow (Fig. 1(b)) that uses a combination of machine learning interatomic potentials (MLIPs), the BEE-NET property predictor, and density functional theory (DFT) calculations to assess stability and predict T_{c}. The following subsections detail each step of this computational process.

2.1 Data

We used two main datasets in this study: the Alexandria Materials Database, and a labeled dataset of superconductors Cerqueira et al. [2024] included the superconducting materials predicted in BEE-NET Gibson et al. [2025b]. The Alexandria Materials Database Schmidt et al. [2024, 2022b, 2023b] contains over 5 million crystal structures. To train our foundation model, we restricted the data to structures with up to 20 atoms per unit cell, giving us 2,086,767 structures, divided into subset of 1,857,222 crystal structures for training and another 229,545 for validation. This large, unlabeled dataset enables the model to learn the manifold of material space. For fine-tuning, we used the DS-A dataset of Cerqueira et al. Cerqueira et al. [2023], comprising 7,217 dynamically stable metallic compounds with first-principles electron–phonon results (relaxed structures, density of states at the Fermi level, logarithmic average phonon frequency, and electron–phonon coupling). We used the Allen–Dynes formula Allen and Dynes [1975] to estimate T_{c} of these materials, and retaining 6,326 superconducting materials, reduction due to the numerical instability from the Allen-Dynes formula. We then augmented this set with 857 additional superconducting candidates from Gibson et al. [2025b], yielding 7,183 labeled samples for fine-tuning. The fine-tuning data were split into training/validation/test in an 8:1:1 ratio, and early stopping was used during training. We use only the T_{c} values for fine-tuning.

2.2 Foundation Model Training

The first stage in our methodology is the development of a robust foundation model for crystal structure generation. The primary objective of this foundational training is to create a model adept at generating a diverse range of structurally and thermodynamically plausible crystalline materials, capturing the inherent chemical and physical rules governing crystal stability.

Training was performed using the DiffCSP framework, the architecture of which is detailed in Section 2.3. During this unsupervised phase, the model learns to reverse a diffusion process that gradually adds noise to the lattice parameters, fractional atomic coordinates, and atom types of known crystal structures. By learning to effectively denoise these inputs, the model implicitly learns the complex manifold of crystal structures. This process is guided by the need to respect fundamental physical symmetries inherent in crystalline materials, such as rotational, translational and periodic boundary conditions. The resulting foundation model is therefore well-versed in the general principles of crystal structure formation and stability, without being biased towards any specific target property, providing a strong starting point for subsequent fine-tuning towards superconductor discovery.

2.3 DiffCSP Architecture

For crystal structure generation within our work, we employ a diffusion model. This generative approach involves a forward process where Gaussian noise is progressively added to the key attributes of a crystal structure - specifically its lattice parameters, the fractional coordinates of its constituent atoms, and their elemental types. The model is then trained to reverse this noising process, learning to denoise these attributes at various noise levels to reconstruct the original, physically valid crystal structure. New crystal structures are subsequently generated by iteratively applying this learned denoising function, beginning from a state of pure noise.

Generating physically realistic crystal structures critically requires the model's adherence to fundamental crystallographic symmetries. These include O(3) equivariance, which ensures that predictions are independent of crystal orientation by requiring consistent transformation of vectorial features (like atomic positions or lattice vectors) under rotation. Additionally, proper handling of periodic boundary conditions (PBC) is essential for accurately representing fractional atomic coordinates and ensuring the continuity of structural properties across unit cell boundaries.

To implement a diffusion process that rigorously incorporates these essential symmetries, we selected the Crystal Structure Prediction by Joint Equivariant Diffusion (DiffCSP) Jiao et al. [2023] framework as the backbone denoising architecture. DiffCSP employs an E(n) equivariant graph neural network (EGNN) Satorras et al. [2021] for the denoising model. This choice ensures O(3)-equivariance for handling lattice and coordinate symmetries and effectively manages periodic translational invariance for fractional atomic coordinates.

The input features consist of atomic embeddings and sinusoidal positional encoding that represent atom types and their positions within the crystal lattice. These features are processed through multiple message-passing layers, which incorporate both spatial and periodic symmetries. A Fourier transform is applied to the relative fractional coordinates, improving the model's ability to capture periodic behavior. The denoising model outputs two components: one for lattice noise and one for fractional coordinate noise. These outputs are designed to respect both O(3)-equivariance and translational invariance. Together, these architectural properties enable the model to robustly denoise inputs Ho et al. [2020], Song et al. [2021] in a way that preserves the physical symmetries of crystal structures.

In addition to lattice and coordinate modeling, atom types are also handled through learned embeddings/denoising, though detailed discussion of atom-type noise is omitted here for brevity. We kept the model hyperparameters consistent with those used in the original DiffCSP framework.

2.4 Fine-tuning for Conditional Generation

The second stage of training involves fine-tuning the foundation model towards generation of superconducting materials. We use an adapter module Zeni et al. [2024] into the equivariant graph network (CSPNet) used in DiffCSP; similar ideas have already been implemented in computer vision for text-to-image generation Rombach et al. [2021], Zhang et al. [2023], Mou et al. [2023]. This module allows us to condition generation on the target property, in this case, scalar T_{c} values. The adapter module is applied after each message-passing layer, at layer L, node embeddings are updated as

$$H_j^{'(L)} = H_j^{(L)} + f_{\text{mixin}}^{(L)}(f_{\text{adapter}}^{(L)}(g)) \cdot I(\text{property is not null})$$

where $H_j^{(L)}$ denotes the embedding of node j at layer L, and g is a sinusoidal embedding Vaswani et al. [2017] of the scalar target T_{c}. The adapter $f_{\text{adapter}}^{(L)}$ consists of a small stack of four fully connected layers. Its output is passed through a mixin layer Harrison and Ossher [2001] $f_{\text{mixin}}^{(L)}$ that gates the adapter signal, stabilizing training by gradually increasing the adapter's influence. The mixin layer initially scales the adapter output to zero and increases it progressively during fine-tuning. This mechanism prevents abrupt disruptions to the foundation model's learned features.

We retain the original DiffCSP denoising objective, so the pretrained structural prior is preserved while the adapters inject T_{c} information. This yields a property-conditioned generator without degrading the foundation model's learned notions of structural plausibility.

2.5 Classifier-free Guided Generation

To generate candidate superconducting materials, we implemented classifier-free diffusion Ho and Salimans [2021] guidance. This method enables conditional generation by interpolating between the unconditional foundation model and the fine-tuned, property-aware model. Specifically, we combine the denoising predictions from both models as follows:

$$\tilde{\epsilon}_\theta(z_\lambda, c) = (1 + w)\epsilon_\theta(z_\lambda, c) - w\epsilon_\theta(z_\lambda)$$

Here, $\epsilon_\theta(z_\lambda, c)$ is the denoising output of the conditional model guided by the target property c, and $\epsilon_\theta(z_\lambda)$ is the output of the unconditional model. The guidance weight w controls the strength of the conditioning — higher values bias the generation more strongly toward the target property. We found $w = 2$ to be a stable and effective choice in practice.

Using this guided generation strategy, we sampled 200,000 crystal structures by conditioning on different T_c values. These generated candidates were then passed through our structure analysis workflow to identify stable, high-T_c superconductors.

2.6 Structural Analysis Workflow

We follow a rigorous multi-step workflow to identify stable, high-T_c superconducting candidates from the set of generated structures. This pipeline ensures that final materials are metallic, thermodynamically and dynamically stable, and synthesizable. Before these generated structures enter the detailed screening pipeline, an initial filtering step is performed to ensure a focus on novel candidates. This involves removing any generated structures that are identical to those present in our fine-tuning dataset. Additionally, duplicate structures within the generated set itself are identified and reduced to unique instances. This pre-screening ensures that the subsequent computationally intensive analyses are concentrated on genuinely new potential superconductors.

First, we relax all remaining unique, novel generated structures using the M3GNet machine learning interatomic potential (MLIP) Chen and Ong [2022]. After relaxation, we use MEGNet Chen et al. [2019] to compute the bandgap (E_g) and formation energy (E_f). We retain only metallic structures ($E_\mathrm{g} = 0$) with negative E_f. Next, we estimate the superconducting critical temperature (T_c) of these filtered structures using BEE-Net Gibson et al. [2025b]. We only keep materials with predicted $T_c > 5$ K for further analysis. To assess thermodynamic stability, we calculate the energy above the convex hull (E_hull) using M3GNet and compare against the Materials Project database Jain et al. [2013]. Structures with $E_\mathrm{hull} > 200$ meV/atom are discarded. We further refine the surviving structures by recalculating E_hull using DFT-relaxed geometries and reapply the $T_c > 5$ K filter using BEE-Net.

At this point, we incorporate phonon calculations to improve both dynamic stability assessment and T_c accuracy. For each structure, we compute the phonon density of states (PhDOS) and use BEE-Net with coarse phonon density (CPD) embeddings to predict a more refined T_c. Structures with imaginary phonon modes are eliminated. In the final stage, we perform electron-phonon coupling calculations using Quantum ESPRESSO to obtain the electron-phonon spectral function $\alpha^2F(\omega)$ and compute the final DFT-based T_c values using the Allen-Dynes equation Allen and Dynes [1975].

3 Results

We first validate our computational approach by demonstrating that guidance allows for the generation of materials with desired superconducting critical temperatures. We then characterize the full set of generated structures with theoretical predictions and find novel superconducting candidates.

3.1 Validation of the Guided Diffusion model

We validated the performance of the guided diffusion model by assessing its ability to generate structures conditioned on a target T_c. The guidance mechanism, based on classifier-free guidance Ho and Salimans [2021], is designed to balance this drive toward a specific T_c with the principles of structural viability by the foundation model (Sec. Foundation model and Sec. Fine Tuning). This balance is tuned by a guidance weight w (detailed in Sec. Classifier-free Guidance). To validate our guided diffusion model we analyze the distributions of T_c values for sets of 1,000 generated structures,

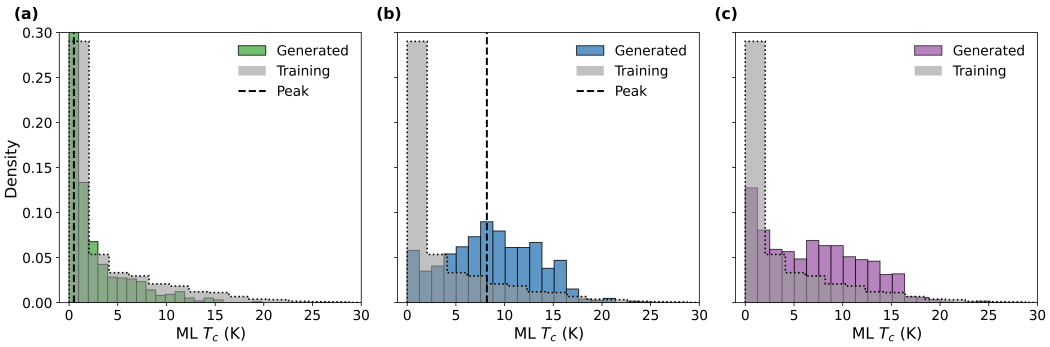

Figure 2: Effect of classifier-free guidance on T_c distribution of 1,000 generated structures per setting, with T_c predicted by BEE-NET with the crystal structure only (CSO) input. (a) Generated structures with guidance disabled ($w = -1$) closely follow the fine-tuning dataset distribution. (b) With guidance enabled ($w = 2$) and target $T_c = 10$ K, the distribution shifts toward the desired value, demonstrating the model's controllability. (c) With $w = 2$ and target $T_c = 110$ K - well outside the training distribution - the sampled distribution remains concentrated at low T_c, indicating adherence to the learned plausible structural prior rather than oversteering to satisfy the property target.

estimating T_c using crystal structure only (CSO) input to BEE-NET Gibson et al. [2025b] We did not perform any structural stability test on the generated candidates for this step. This approach provides near-immediate feedback without the computational cost of ab initio calculations, which are reserved for the more comprehensive analysis workflow described in Sec. Structural Analysis Workflow. As to the efficacy of DiffCSP to generate stable and novel structures, we refer the reader to Ref. Jiao et al. [2023].

Our analysis illustrated in Fig. 2 first confirms that with guidance disabled ($w = -1$), the T_c distribution of the generated samples closely follows that of the fine-tuning dataset, establishing a baseline (Fig. 2a). In contrast, enabling classifier-free guidance (e.g., $w = 2$) to target a specific T_c of 10 K results in a clear and effective shift of the distribution towards the desired value, demonstrating property-driven control (Fig. 2b). We further probed the model behavior with an out-of-distribution target of $T_c = 110$ K, a regime absent from the fine-tuning data. Even with strong guidance, the model consistently failed to produce high-T_c structures. Instead, as shown in Fig. 2c ($w = 2$), the model preferentially generates low-T_c structures. This behavior demonstrates that while guidance is effective for targeted design within the training domain, the foundation model acts as a crucial prior for stability, constraining the generation to physically plausible structures for extreme property extrapolation.

3.2 Prediction of Superconducting Candidates

We applied the full generative–screening pipeline to 200,000 initial structures. After fast MEGNet Chen et al. [2019] relaxation and deduplication, 34,027 unique and novel generated structures remained. These were subjected to the multistage structural-analysis workflow (detailed in subsection Structural Analysis Workflow).

Of the 34,027 identified structures, 20,994 were classified as metals with negative formation energy ($E_f < 0$). From this set, 10,877 were predicted to have $T_c > 5$ K by BEE-Net using crystal structure only (CSO) inputs. Thermodynamic pre-screening with M3GNet retained 9,167 with energy above the convex hull $E_{\text{hull}} < 200$ meV/atom. After DFT relaxation, 8,014 structures converged; 6,173 of these satisfied a DFT-calculated $E_{\text{hull}} < 200$ meV/atom Sun et al. [2016]. Phonon calculations identified 1,818 as dynamically stable. Incorporating coarse phonon-density (CPD) embeddings into BEE-Net yielded 1,365 with predicted $T_c > 5$ K, from which 773 unique candidates exhibit DFT-calculated $T_c > 5$ K. The relationship between predicted T_c and E_{hull} for these candidates is shown in Figure 3, spanning T_c values up to $\sim$35 K and highlighting materials that combine promising superconducting properties with thermodynamic stability conducive to synthesis.

A compositional analysis of these 773 candidates revealed a strong trend towards multi-component compounds with 133 binaries (17%), 455 ternaries (59%), 178 quaternaries (23%), and 7 pentanaries

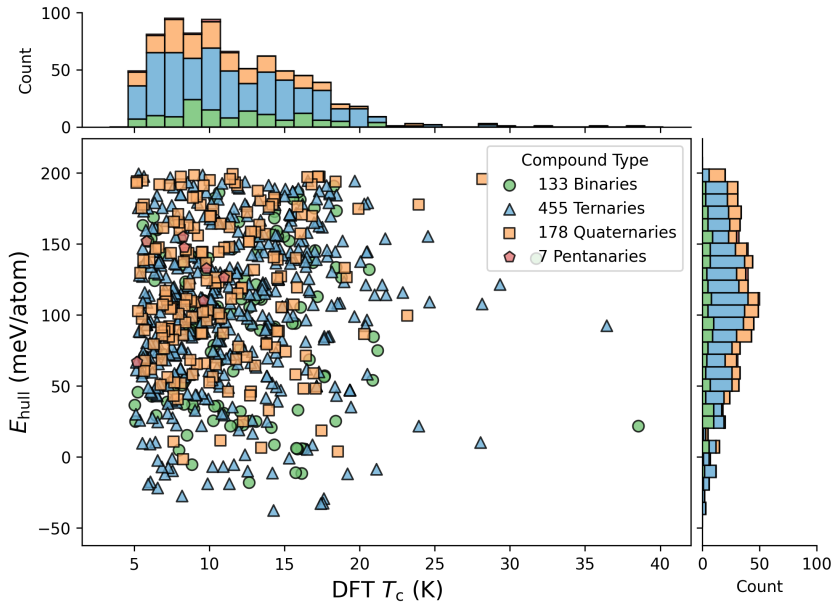

Figure 3: Distribution of the 773 predicted superconducting candidates. The main panel shows energy above the convex hull (E_{hull}) versus DFT-calculated T_{c}; marker color/shape denotes compound type as indicated in inset legend. Marginal histograms along the top (for T_{c}) and right (for E_{hull}) display stacked counts by compound type, using 30 equal-width bins aligned to the main panel axis limits.

(1%). We attribute this trend to the combined effect of our model's generative capabilities and the current state of superconductor research. The landscape of binary superconductors has been extensively investigated. In essence, for the workflow to identify new and viable superconductors, it is more probable that these will emerge from these more complex systems where a greater potential for undiscovered materials exists, rather than from the more saturated binary space.

Our generative AI approach proved to be highly effective, as out of 34,027 unique structures 773 are predicted to be superconductors with $T_{\text{c}} > 5K$. For comparison, an element-substitution baseline approach using the same BEE-Net/DFT workflow identified 204 superconductors from 1.22 million unique candidates (0.017%), whereas the here-proposed generative pipeline yields 773 of 34,027 (2.3%), a $\sim$135$\times$ improvement in hit rate Gibson et al. [2025b]. Within generative approaches, Wines et al. generated 3,000 candidates and obtained 25 materials with DFT-calculated $T_{\text{c}} > 5$ K (0.83%) Wines et al. [2023]; the present pipeline achieves a higher per-structure hit rate (2.3% vs. 0.83%, $\sim$2.8$\times$) and a larger absolute yield (773 vs. 25, $\sim$31$\times$).

4 Discussion

In this work, we developed and deployed an end-to-end workflow that combines a guided diffusion generative model with high-throughput computational screening to accelerate the inverse design of novel superconductors. As quantified in the Results, our pipeline achieves a 2.3% per-structure hit rate (773/34,027), improving on an element-substitution baseline (0.017%, 204/1.22 million Gibson et al. [2025b]) by $\sim$135$\times$; relative to a recent generative search (25/3,000 Wines et al. [2023]), it also delivers a $\sim$31$\times$ larger absolute yield (773 vs. 25).

The distribution of stability, i.e. E_{hull}, as shown in Fig.3, can likely be attributed to the pre-training and fine-tuning datasets, both consist mostly of theoretical materials that are not all on the convex hull. It is likely that training on experimentally synthesized materials, and/or only materials with E_{hull} near zero, would guide the model to generate candidates that are, on average, closer to the hull.

As noted in the Results section, the final candidate set skews toward ternary compositions, reflecting the extensive prior exploration of binaries. Many of these ternaries involve relatively costly elements, suggesting historical under-sampling of such chemistries in experiments. Future iterations should

encode cost and handling constraints (e.g., excluding toxic or radioactive elements) as generative priors and screening filters, steering discovery toward candidates that are both scientifically promising and experimentally tractable. At the same time, pushing into underexplored chemistries raises a distinct caveat we discuss below: data sparsity can make phases appear stable relative to an incomplete convex hull.

A critical limitation arises in underexplored chemistries. The model's prediction of stable phases in these regions could be an artifact of data scarcity, not an indication of true physical plausibility. For a predicted structure to be deemed stable, its energy must be low relative to a convex hull constructed from known, competing phases. In chemical spaces where data is sparse – for instance, among ternary systems whose constituent elements do not readily form any known binary compounds – the reference convex hull is likely incomplete. Consequently, a generated structure may appear to have a low energy above the hull (E_{hull}) simply because there are no known, more stable structures to compete with it in the database. Our generative model, therefore, may be expertly identifying gaps in existing materials data rather than discovering genuinely synthesizable, stable phases. Indeed, close to two thirds of the binary phase diagrams that are relevant to the materials predicted by the model contain no known ordered binary compounds, often because the elements form complete solid solutions.

The factors above highlight a fundamental challenge for AI-driven materials discovery: distinguishing true, synthesizable novelty from artifacts of an incomplete reference dataset. Therefore, while our work validates the immense potential of generative AI to accelerate materials discovery, it also underscores the critical need for future developments to focus on improving predictions of synthesizability and integrating experimental feedback into a true active learning loop to bridge the gap between computational prediction and laboratory realization.

5 Conclusion

We have demonstrated that guided diffusion models can significantly accelerate the discovery of novel superconductor candidates, achieving a 135$\times$ improvement in hit rate over element-substitution approaches. However, critical challenges remain in synthesizability assessment, particularly in underexplored chemical spaces where data sparsity can create artifacts in stability predictions.

Future directions could involve improvements to the generative model itself. Replacing fixed property embeddings with learnable ones and evaluating alternative guidance schemes Karras et al. [2024], Tang et al. [2025] may improve generalization and coverage relative to classifier-free guidance. To better reflect the discrete nature of the composition, guided discrete diffusion for element identities Nisonoff et al. [2025] may be preferable over denoising in a continuous latent space. It is also natural to integrate our guidance framework with alternative generative backbones—including Flow Matching Miller et al. [2024] and stochastic interpolants Höllmer et al. [2025]—to assess their efficacy in discovering novel, synthesizable superconductors.

Bridging prediction and synthesizability points to several clear next steps. To address sparse reference data in undercharted chemical spaces, on-the-fly DFT refinement of local convex hulls around promising candidates could be employed. An experiment-in-the-loop active-learning workflow—incorporating non-formations as hard negatives, retraining the model iteratively, and expanding training data with experimentally verified compounds—could sharpen decision boundaries. Multi-objective optimization that balances high T_{c} with synthesizability and practical constraints (cost, handling, toxicity) could further focus searches. Finally, disorder-aware screening that compares predicted ordered compounds against competing solid solutions could improve experimental realizability and hit rates.

References

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, volume 33, pages 6840–6851. Curran Associates, Inc., 2020. URL `https://proceedings.neurips.cc/paper/2020/hash/4c5bcfec8584af0d967f1ab10179ca4b-Abstract.html`.

Rui Jiao, Wenbing Huang, Peijia Lin, Jiaqi Han, Pin Chen, Yutong Lu, and Yang Liu. Crystal structure prediction by joint equivariant diffusion. In *Thirty-seventh Conference on Neural Information*

Processing Systems, 2023. URL https://openreview.net/forum?id=DNdN26m2Jk.

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

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.

Benjamin Kurt Miller, Ricky T. Q. Chen, Anuroop Sriram, and Brandon M Wood. FlowMM: Generating materials with riemannian flow matching. In *Forty-first International Conference on Machine Learning*, 2024. URL https://openreview.net/forum?id=W4pB7VbzZI.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30, 2017. URL https://arxiv.org/abs/1706.03762.

Luis M. Antunes, Keith T. Butler, and Ricardo Grau-Crespo. Crystal structure generation with autoregressive large language modeling. *Nature Communications*, 15(1), December 2024. ISSN 2041-1723. doi: 10.1038/s41467-024-54639-7. URL http://dx.doi.org/10.1038/s41467-024-54639-7.

Pierre-Paul De Breuck, Hashim A. Piracha, Gian-Marco Rignanese, and Miguel A. L. Marques. A generative material transformer using wyckoff representation, 2025. URL https://arxiv.org/abs/2501.16051.

Michael S. Albergo, Nicholas M. Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions, 2023. URL https://arxiv.org/abs/2303.08797.

Philipp Höllmer, Thomas Egg, Maya Martirossyan, Eric Fuemmeler, Zeren Shui, Amit Gupta, Pawan Prakash, Adrian Roitberg, Mingjie Liu, George Karypis, Mark Transtrum, Richard Hennig, Ellad B. Tadmor, and Stefano Martiniani. Open materials generation with stochastic interpolants. In *Forty-second International Conference on Machine Learning*, 2025. URL https://openreview.net/forum?id=gHGrzxFujU.

Lilia Boeri, Richard Hennig, Peter Hirschfeld, Gianni Profeta, Antonio Sanna, Eva Zurek, Warren E Pickett, Maximilian Amsler, Ranga Dias, Mikhail I Eremets, Christoph Heil, Russell J Hemley, Hanyu Liu, Yanming Ma, Carlo Pierleoni, Aleksey N Kolmogorov, Nikita Rybin, Dmitry Novoselov, Vladimir Anisimov, Artem R Oganov, Chris J Pickard, Tiange Bi, Ryotaro Arita, Ion Errea, Camilla Pellegrini, Ryan Requist, E K U Gross, Elena Roxana Margine, Stephen R Xie, Yundi Quan, Ajinkya Hire, Laura Fanfarillo, G R Stewart, J J Hamlin, Valentin Stanev, Renato S Gonnelli, Erik Piatti, Davide Romanin, Dario Daghero, and Roser Valenti. The 2021 room-temperature superconductivity roadmap. *Journal of Physics: Condensed Matter*, 34(18):183002, mar 2022. doi: 10.1088/1361-648X/ac2864. URL https://dx.doi.org/10.1088/1361-648X/ac2864.

David Larbalestier, Alex Gurevich, D. Matthew Feldmann, and Anatoly Polyanskii. High-T_c superconducting materials for electric power applications. *Nature*, 414(6861):368–377, 2001. doi: 10.1038/35104654.

Alexis P. Malozemoff, Jochen Mannhart, and Douglas Scalapino. High-temperature cuprate superconductors get to work. *Physics Today*, 58(4):41–47, 04 2005. ISSN 0031-9228. doi: 10.1063/1.1955478. URL https://doi.org/10.1063/1.1955478.

Rene Flükiger. Overview of superconductivity and challenges in applications. *Reviews of Accelerator Science and Technology*, 05:1–23, 2012. doi: 10.1142/S1793626812300010. URL https://doi.org/10.1142/S1793626812300010.

J. Bardeen, L. N. Cooper, and J. R. Schrieffer. Theory of superconductivity. *Phys. Rev.*, 108:1175–1204, Dec 1957. doi: 10.1103/PhysRev.108.1175. URL https://link.aps.org/doi/10.1103/PhysRev.108.1175.

G M Eliashberg. Interactions between electrons and lattice vibrations in a superconductor. *Sov. Phys. - JETP (Engl. Transl.); (United States)*, 9 1960. URL https://www.osti.gov/biblio/7354388.

Paolo Giannozzi, Oscar Baseggio, Pietro Bonfà, Davide Brunato, Roberto Car, Ivan Carnimeo, Carlo Cavazzoni, Stefano de Gironcoli, Pietro Delugas, Fabrizio Ferrari Ruffino, Andrea Ferretti, Nicola Marzari, Iurii Timrov, Andrea Urru, and Stefano Baroni. Quantum espresso toward the exascale. *The Journal of Chemical Physics*, 152(15):154105, 04 2020. ISSN 0021-9606. doi: 10.1063/5.0005082. URL https://doi.org/10.1063/5.0005082.

S. Poncé, E.R. Margine, C. Verdi, and F. Giustino. Epw: Electron–phonon coupling, transport and superconducting properties using maximally localized wannier functions. *Computer Physics Communications*, 209:116–133, December 2016. ISSN 0010-4655. doi: 10.1016/j.cpc.2016.07.028. URL http://dx.doi.org/10.1016/j.cpc.2016.07.028.

S. R. Xie, G. R. Stewart, J. J. Hamlin, P. J. Hirschfeld, and R. G. Hennig. Functional form of the superconducting critical temperature from machine learning. *Physical Review B*, 100(17), November 2019. ISSN 2469-9969. doi: 10.1103/physrevb.100.174513. URL http://dx.doi.org/10.1103/PhysRevB.100.174513.

S. R. Xie, Y. Quan, A. C. Hire, B. Deng, J. M. DeStefano, I. Salinas, U. S. Shah, L. Fanfarillo, J. Lim, J. Kim, G. R. Stewart, J. J. Hamlin, P. J. Hirschfeld, and R. G. Hennig. Machine learning of superconducting critical temperature from eliashberg theory. *npj Computational Materials*, 8(1), January 2022a. ISSN 2057-3960. doi: 10.1038/s41524-021-00666-7. URL http://dx.doi.org/10.1038/s41524-021-00666-7.

Jason B. Gibson, Ajinkya C. Hire, Philip M. Dee, Oscar Barrera, Benjamin Geisler, Peter J. Hirschfeld, and Richard G. Hennig. Accelerating superconductor discovery through tempered deep learning of the electron-phonon spectral function. *npj Computational Materials*, 11(1), January 2025a. ISSN 2057-3960. doi: 10.1038/s41524-024-01475-4. URL http://dx.doi.org/10.1038/s41524-024-01475-4.

Jason B. Gibson, Ajinkya C. Hire, Philip M. Dee, Benjamin Geisler, Jung Soo Kim, Zhongwei Li, James J. Hamlin, Gregory R. Stewart, P. J. Hirschfeld, and Richard G. Hennig. Developing a complete ai-accelerated workflow for superconductor discovery, 2025b. URL https://arxiv.org/abs/2503.20005.

Tiago F. T. Cerqueira, Antonio Sanna, and Miguel A. L. Marques. Sampling the materials space for conventional superconducting compounds. *Advanced Materials*, 36(1), December 2023. ISSN 1521-4095. doi: 10.1002/adma.202307085. URL http://dx.doi.org/10.1002/adma.202307085.

Daniel Wines, Tian Xie, and Kamal Choudhary. Inverse design of next-generation superconductors using data-driven deep generative models. *The Journal of Physical Chemistry Letters*, 14(29): 6630–6638, July 2023. ISSN 1948-7185. doi: 10.1021/acs.jpclett.3c01260. URL http://dx.doi.org/10.1021/acs.jpclett.3c01260.

Tian Xie, Xiang Fu, Octavian-Eugen Ganea, Regina Barzilay, and Tommi S. Jaakkola. Crystal diffusion variational autoencoder for periodic material generation. In *International Conference on Learning Representations*, 2022b. URL https://openreview.net/forum?id=03RLpj-tc_.

Kamal Choudhary and Brian DeCost. Atomistic line graph neural network for improved materials property predictions. *npj Computational Materials*, 7(1):185, 2021. doi: 10.1038/s41524-021-00650-1. URL https://www.nature.com/articles/s41524-021-00650-1.

Samuel Yuan and S. V. Dordevic. Diffusion models for conditional generation of hypothetical new families of superconductors. *Scientific Reports*, 14(1):10275, 2024. doi: 10.1038/s41598-024-61040-3. URL https://www.nature.com/articles/s41598-024-61040-3.

Jooyoung Choi, Sungwon Kim, Yonghyun Jeong, Youngjune Gwon, and Sungroh Yoon. Ilvr: Conditioning method for denoising diffusion probabilistic models. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14347–14356, 2021. doi: 10.1109/ICCV48922.2021.01410.

Elizabeth A. Pogue, Alexander New, Kyle McElroy, Nam Q. Le, Michael J. Pekala, Ian McCue, Eddie Gienger, Janna Domenico, Elizabeth Hedrick, Tyrel M. McQueen, Brandon Wilfong, Christine D. Piatko, Christopher R. Ratto, Andrew Lennon, Christine Chung, Timothy Montalbano, Gregory Bassen, and Christopher D. Stiles. Closed-loop superconducting materials discovery. *npj Computational Materials*, 9(1), October 2023. ISSN 2057-3960. doi: 10.1038/s41524-023-01131-3. URL http://dx.doi.org/10.1038/s41524-023-01131-3.

Brandon Wilfong, Alexander New, Gregory Bassen, Wyatt Bustine, Michael J. Pekala, Nam Q. Le, Karun K. Rao, Elizabeth A. Pogue, Eddie Gienger, Michał J. Winiarski, Tyrel M. McQueen, and Christopher D. Stiles. Ternary materials discovery using human-in-the-loop generative machine learning. *RSC Advances*, 15(24):19126–19131, 2025. ISSN 2046-2069. doi: 10.1039/d5ra00427f. URL http://dx.doi.org/10.1039/D5RA00427F.

Michael J. Hutcheon, Alice M. Shipley, and Richard J. Needs. Predicting novel superconducting hydrides using machine learning approaches. *Physical Review B*, 101(14), April 2020. ISSN 2469-9969. doi: 10.1103/physrevb.101.144505. URL http://dx.doi.org/10.1103/PhysRevB.101.144505.

Zhantao Chen, Nina Andrejevic, Tess Smidt, Zhiwei Ding, Qian Xu, Yen-Ting Chi, Quynh T. Nguyen, Ahmet Alatas, Jing Kong, and Mingda Li. Direct prediction of phonon density of states with euclidean neural networks. *Advanced Science*, 8(12), March 2021. ISSN 2198-3844. doi: 10.1002/advs.202004214. URL http://dx.doi.org/10.1002/advs.202004214.

Sean D. Griesemer, Bianca Baldassarri, Ruijie Zhu, Jiahong Shen, Koushik Pal, Cheol Woo Park, and Chris Wolverton. Wide-ranging predictions of new stable compounds powered by recommendation engines. *Science Advances*, 11(1), January 2025. ISSN 2375-2548. doi: 10.1126/sciadv.adq1431. URL http://dx.doi.org/10.1126/sciadv.adq1431.

Benediktus Madika, Aditi Saha, Chaeyul Kang, Batzorig Buyantogtokh, Joshua Agar, Chris M. Wolverton, Peter Voorhees, Peter Littlewood, Sergei Kalinin, and Seungbum Hong. Artificial intelligence for materials discovery, development, and optimization. *ACS Nano*, 19 (30):27116–27158, July 2025. ISSN 1936-086X. doi: 10.1021/acsnano.5c04200. URL http://dx.doi.org/10.1021/acsnano.5c04200.

Jonathan Schmidt, Noah Hoffmann, Hai-Chen Wang, Pedro Borlido, Pedro J. M. A. Carriço, Tiago F. T. Cerqueira, Silvana Botti, and Miguel A. L. Marques. Machine-learning-assisted determination of the global zero-temperature phase diagram of materials. *Advanced Materials*, 35(22), April 2023a. ISSN 1521-4095. doi: 10.1002/adma.202210788. URL http://dx.doi.org/10.1002/adma.202210788.

Jonathan Schmidt, Hai-Chen Wang, Tiago F. T. Cerqueira, Silvana Botti, and Miguel A. L. Marques. A dataset of 175k stable and metastable materials calculated with the pbesol and scan functionals. *Scientific Data*, 9(1), March 2022a. ISSN 2052-4463. doi: 10.1038/s41597-022-01177-w. URL http://dx.doi.org/10.1038/s41597-022-01177-w.

Chi Chen, Weike Ye, Yunxing Zuo, Chen Zheng, and Shyue Ping Ong. Graph networks as a universal machine learning framework for molecules and crystals. *Chemistry of Materials*, 31(9):3564–3572, 2019. doi: 10.1021/acs.chemmater.9b01294. URL https://doi.org/10.1021/acs.chemmater.9b01294.

Chi Chen and Shyue Ping Ong. A universal graph deep learning interatomic potential for the periodic table. *Nature Computational Science*, 2:718–728, 2022. doi: 10.1038/s43588-022-00349-3. URL https://www.nature.com/articles/s43588-022-00349-3.

Tiago F. T. Cerqueira, Antonio Sanna, and Miguel A. L. Marques. Sampling the materials space for conventional superconducting compounds. *Advanced Materials*, 36(1):2307085, 2024. doi: https://doi.org/10.1002/adma.202307085. URL https://advanced.onlinelibrary.wiley.com/doi/abs/10.1002/adma.202307085.

Jonathan Schmidt, Tiago F.T. Cerqueira, Aldo H. Romero, Antoine Loew, Fabian JÃ€ger, Hai-Chen Wang, Silvana Botti, and Miguel A.L. Marques. Improving machine-learning models in materials science through large datasets. *Materials Today Physics*, 48:101560, 2024. ISSN 2542-5293. doi: https://doi.org/10.1016/j.mtphys.2024.101560. URL https://www.sciencedirect.com/science/article/pii/S2542529324002360.

Jonathan Schmidt, Hai-Chen Wang, Tiago F. T. Cerqueira, Silvana Botti, and Miguel A. L. Marques. A dataset of 175k stable and metastable materials calculated with the pbesol and scan functionals. *Scientific Data*, 9(1), March 2022b. ISSN 2052-4463. doi: 10.1038/s41597-022-01177-w. URL http://dx.doi.org/10.1038/s41597-022-01177-w.

Jonathan Schmidt, Noah Hoffmann, Hai-Chen Wang, Pedro Borlido, Pedro J. M. A. Carriço, Tiago F. T. Cerqueira, Silvana Botti, and Miguel A. L. Marques. Machine-learning-assisted determination of the global zero-temperature phase diagram of materials. *Advanced Materials*, 35(22), April 2023b. ISSN 1521-4095. doi: 10.1002/adma.202210788. URL http://dx.doi.org/10.1002/adma.202210788.

P. B. Allen and R. C. Dynes. Transition temperature of strong-coupled superconductors reanalyzed. *Phys. Rev. B*, 12:905–922, Aug 1975. doi: 10.1103/PhysRevB.12.905. URL https://link.aps.org/doi/10.1103/PhysRevB.12.905.

Victor Garcia Satorras, Emiel Hoogeboom, and Max Welling. E(n) equivariant graph neural networks. *CoRR*, abs/2102.09844, 2021. URL https://arxiv.org/abs/2102.09844.

Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations, 2021. URL https://arxiv.org/abs/2011.13456.

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2021. URL https://arxiv.org/abs/2112.10752.

Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models, 2023. URL https://arxiv.org/abs/2302.05543.

Chong Mou, Xintao Wang, Liangbin Xie, Yanze Wu, Jian Zhang, Zhongang Qi, Ying Shan, and Xiaohu Qie. T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models, 2023. URL https://arxiv.org/abs/2302.08453.

William Harrison and Harold Ossher. Mixin layers: An object-oriented implementation technique for refinements and collaboration-based designs. *ACM Transactions on Software Engineering and Methodology (TOSEM)*, 11(2):205–255, 2001.

Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2021. URL https://openreview.net/forum?id=qw8AKxfYbI.

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):011002, 2013. doi: 10.1063/1.4812323.

Wenhao Sun, Stephen T. Dacek, Shyue Ping Ong, Geoffroy Hautier, Anubhav Jain, William D. Richards, Anthony C. Gamst, Kristin A. Persson, and Gerbrand Ceder. The thermodynamic scale of inorganic crystalline metastability. *Science Advances*, 2(11), November 2016. ISSN 2375-2548. doi: 10.1126/sciadv.1600225. URL http://dx.doi.org/10.1126/sciadv.1600225.

Tero Karras, Miika Aittala, Tuomas Kynkäänniemi, Jaakko Lehtinen, Timo Aila, and Samuli Laine. Guiding a diffusion model with a bad version of itself. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL https://openreview.net/forum?id=bg6fVPVs3s.

Zhicong Tang, Jianmin Bao, Dong Chen, and Baining Guo. Diffusion models without classifier-free guidance, 2025. URL https://arxiv.org/abs/2502.12154.

Hunter Nisonoff, Junhao Xiong, Stephan Allenspach, and Jennifer Listgarten. Unlocking guidance for discrete state-space diffusion and flow models. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=XsgHl54yO7.